



Control of movement vigor and decision making during foraging

Tehrim Yoon^a, Robert B. Geary^a, Alaa A. Ahmed^{b,c}, and Reza Shadmehr^{a,1}

^aLaboratory for Computational Motor Control, Department of Biomedical Engineering, Johns Hopkins School of Medicine, Baltimore, MD 21205;

^bDepartment of Integrative Physiology, University of Colorado, Boulder, CO 80309; and ^cDepartment of Mechanical Engineering, University of Colorado, Boulder, CO 80309

Edited by Peter L. Strick, University of Pittsburgh, Pittsburgh, PA, and approved September 25, 2018 (received for review July 31, 2018)

During foraging, animals decide how long to stay at a patch and harvest reward, and then, they move with certain vigor to another location. How does the brain decide when to leave, and how does it determine the speed of the ensuing movement? Here, we considered the possibility that both the decision-making and the motor control problems aimed to maximize a single normative utility: the sum of all rewards acquired minus all efforts expended divided by total time. This optimization could be achieved if the brain compared a local measure of utility with its history. To test the theory, we examined behavior of people as they gazed at images: they chose how long to look at the image (harvesting information) and then moved their eyes to another image, controlling saccade speed. We varied reward via image content and effort via image eccentricity, and then, we measured how these changes affected decision making (gaze duration) and motor control (saccade speed). After a history of low rewards, people increased gaze duration and decreased saccade speed. In anticipation of future effort, they lowered saccade speed and increased gaze duration. After a history of high effort, they elevated their saccade speed and increased gaze duration. Therefore, the theory presented a principled way with which the brain may control two aspects of behavior: movement speed and harvest duration. Our experiments confirmed many (but not all) of the predictions, suggesting that harvest duration and movement speed, fundamental aspects of behavior during foraging, may be governed by a shared principle of control.

decision making | motor control | foraging | vigor | theoretical neuroscience

Animals make foraging decisions that, from a local perspective, may seem illogical. For example, a crow may spend effort digging up a small clam from a sandy beach but then abandon it in favor of searching for a larger clam (1). This behavior is locally suboptimal; why not open the clam and eat it before moving on? To understand these decisions, ecologists have considered a global utility called optimal foraging (2). In this utility, decisions aim to maximize the sum of all rewards acquired minus all efforts expended over total time (i.e., the global capture rate). Abandoning food is worthwhile if the time required to consume it can be better used exploiting other options.

Optimal foraging is attractive for two reasons. From a theoretical perspective, it relies on a normative framework, while from an evolutionary perspective, it maximizes a utility that covaries with fecundity and longevity (3). This has made optimal foraging a powerful theory with which to understand decision making in diverse organisms from worms (4) to birds (5, 6) and primates (7–10). For example, the theory correctly predicts that people will abandon an apple tree when the rate of harvest falls below the average rate available in the environment (9). It correctly predicts that monkeys will stay longer and harvest a diminishing juice supply when the wait time to the next opportunity has increased (7).

However, optimal foraging is generally not concerned with the question of movement vigor. This omission is noteworthy, because on the one hand, movement speed affects energetic expenditure (11–13) (one of the key variables in the global capture

rate) and on the other hand, factors that affect decision making (such as reward and effort) also affect motor control (14–19). The first attempt to address this limitation was by Niv et al. (20) who defined a global utility along the same principles used in optimal foraging but with a vigor-dependent effort cost. Recent experiments have shown that patterns of reaction time (21) and decision making (9) are largely consistent with predictions of the theory. However, control of vigor (i.e., speed of movement) in the framework of optimal foraging has remained an open question (22).

Here, we solved the general theoretical problem: given the objective of maximizing the global capture rate, how long should one stay and harvest reward at the current patch, and then, how fast should one travel to the next patch? Our solution relied on marginal value theorem (MVT) (23), which states that, to maximize the global utility, the animal should compare the local capture rate in the current patch (reward acquired minus effort expended divided by time at current location) with the average rate available in the environment, the global capture rate. The results predicted that the interaction between these two rates should determine how fast the subject moved as well as how much time she spent at her destination.

To test the theory, we performed foraging-like experiments where reward patches consisted of small images. Subjects harvested reward by gazing at the image and then moved their eyes with a saccade to view another image. We manipulated reward magnitude via image content and effort expenditure via image eccentricity. We found that people responded to increases in the local availability of reward (improved image content) by increasing their gaze duration. In contrast, after a history of high-valued

Significance

How long should one stay and accumulate reward, and then, how fast should one travel to the next reward site? Marginal value theorem describes the decision-making process: the brain compares the immediate rate of harvest with its global history of capture, deciding to leave when the immediate rate falls below the average. Here, we extended the theory, showing that the same principle can be used to control speed of movements: the brain should compare the immediate rate of energy expenditure during movement with the global capture rate, planning to arrive at the destination when the two become equal. Experimental results confirmed many of the predictions, suggesting that a shared principle may underlie decision making and control of movement vigor.

Author contributions: T.Y., R.B.G., A.A.A., and R.S. designed research; T.Y. and R.B.G. performed research; A.A.A. and R.S. contributed new reagents/analytic tools; T.Y. and R.B.G. analyzed data; and R.S. wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission.

Published under the PNAS license.

¹To whom correspondence should be addressed. Email: shadmehr@jhu.edu.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1812979115/-DCSupplemental.

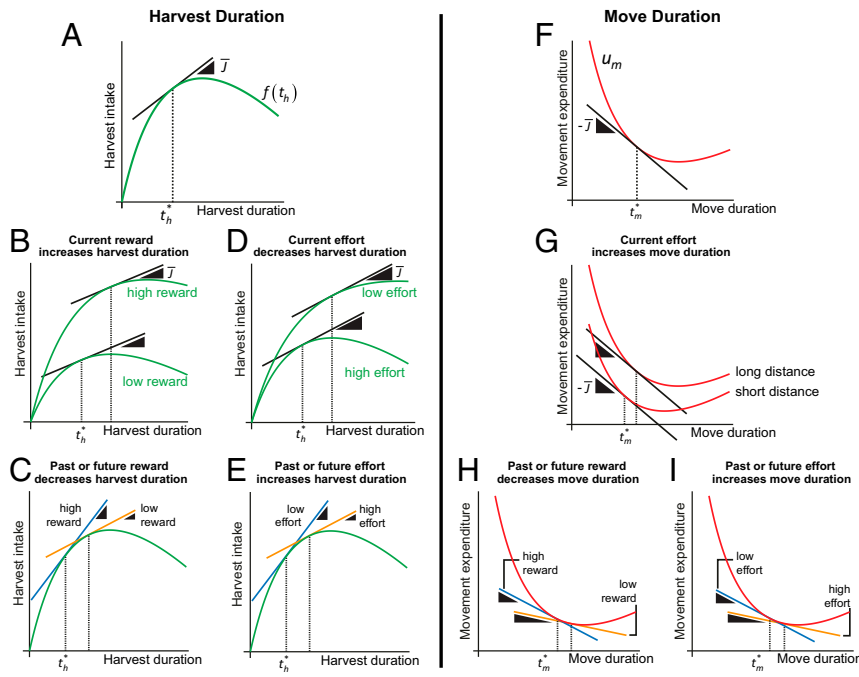


Fig. 1. The theoretical link between harvest duration and movement vigor during foraging. (A) During harvest, reward is accumulated, while effort is expended (Eq. 1), resulting in harvest intake $f(t_h)$, a function that is concave downward. The patch should be abandoned not when the intake is maximized but when the rate of intake (local capture rate) is equal to the average capture rate \bar{J} (i.e., at time t_h^*). (B) Increased reward at the current location extends the harvest duration. Subject should stay longer when there is greater reward. (C) After experience of high reward, indicating a rich environment, subject should stay a shorter period in the current patch. Similarly, in anticipation of high reward in a future harvest, subject should shorten harvest duration at the current patch. (D) Harvest duration should be shortened when the current harvest requires a large amount of effort. (E) Harvest duration should be lengthened if the past history has included high effort. Similarly, if future harvests are expected to encounter high effort, current harvest should be extended. (F) During movement, energy expenditure is a concave upward function of duration. Movement duration is optimal not when the effort cost of travel u_m is minimized but when the rate of energetic loss during travel is equal to the negative of the average capture rate. (G) Increased effort requirements of travel result in increased movement duration. (H) After experience of high reward (larger J) or in anticipation of high reward in the future, movement duration t_m^* should be decreased. That is, vigor should be increased after experience of high reward or while moving toward a high reward patch. (I) After experience of high effort (lower J) or in anticipation of high effort, movement duration t_m^* should be increased. That is, vigor should be reduced after experience of high effort.

rewards, they spent less time at the current image. Saccade vigor increased with expectations of future reward and also increased after a history of high rewards. Subjects decreased gaze duration when the current image required high effort (image eccentricity), but after expenditure of high effort in viewing past images, they spent greater time gazing at the current image. All of these results were predicted by the theory. However, in one respect, people behaved differently than predicted: after experience of high effort, they consistently chose to increase their vigor.

Our results suggest a principled way with which the brain may control both the decision regarding how long to stay and harvest reward and the vigor of the ensuing movement: during foraging, both seem to be controlled via a comparison between the capture rate available at the local environment, with a global rate that depends on the history and expectations of future.

Results

Rewards are often distributed in patches. We move to a patch, stay for some period of time to harvest a portion of the reward, and then move to another patch to acquire another reward. Effort is expended during the movement as well as during the harvest period. Here, our concern is with regard to two questions. (i) How fast should we move from one site to another (the motor control problem)? (ii) How long should we stay at our destination (the decision-making problem)? We will consider these questions in the normative framework of optimal foraging, where the objective is to maximize the total reward harvested across all patches minus the total effort expended divided by total time (i.e., the global capture rate). We provide a theoretical

solution to this problem and then test the predictions of the theory in a series of experiments.

Generalized MVT. There are patches $n = 1, \dots, N$ that are distributed throughout a field, with each patch containing reward quantity $\alpha^{(n)}$. Given the goal of traveling to patch n , the subject can control two variables: duration of time spent traveling to that patch, represented by $t_m^{(n)}$ (movement period), and duration of time spent harvesting reward at that patch, represented by $t_h^{(n)}$ (harvest period).

In our foraging-like experiments, the reward patch is a small image placed somewhere on a screen. The subject spends time gazing at the image, akin to harvesting some portion of the reward. We control reward quantity α via image type and effort required for harvesting via image eccentricity. We are interested in predicting the duration of time that the subject maintains her gaze at that image and the velocity with which she moves her eyes to the next image.

During the harvest period, reward is accumulated via a function that depends on the total amount $\alpha^{(n)}$ available at the patch (image content) and the rate β with which this reward can be harvested. In addition, the act of harvesting requires effort (in this case, gaze holding at some eccentric location), which we represent via $u_h(t_h)$. Let us assume that reward accumulation is an increasing function that rises rapidly but then saturates, indicative of scenarios where reward is finite and spending greater time has diminishing returns. The sum of reward acquired and effort expended at the patch is represented via the harvest function (Fig. 1A):

$$f^{(n)}(t_h) = \alpha^{(n)} \left(1 - \frac{1}{1 + \beta t_h} \right) - u_h^{(n)}(t_h). \quad [1]$$

The exact form of this function is not critical here. What is important is that the harvest function is initially increasing, with a second derivative that is always negative (i.e., concave downward) (Fig. 1A). At patch n , effort expenditure depends on harvest duration as well as an effort rate parameter $k^{(n)}$:

$$u_h^{(n)}(t_h) = k^{(n)} t_h. \quad [2]$$

For example, a more eccentric location of the image will require greater effort expenditure [i.e., a larger $k^{(n)}$]. The harvest function is initially increasing and concave downward if α , β , and k are all positive quantities. By manipulating image content and its location, we will modulate the harvest function.

To acquire reward, one must travel to the patch. During the travel period, effort is expended via a function that depends on distance $d^{(n)}$ and duration of travel $u_m(d^{(n)}, t_m^{(n)})$. To consider the events during travel and harvesting together, optimal foraging suggests that a relevant measure is the local capture rate $J^{(n)}$, defined as the sum of reward accumulated and effort expended at patch n , divided by time that it took to travel and harvest reward:

$$J^{(n)} = \frac{f^{(n)}(t_h^{(n)}) - u_m(d^{(n)}, t_m^{(n)})}{t_h^{(n)} + t_m^{(n)}}. \quad [3]$$

Given that we have many reward sites, the objective is to maximize the global capture rate \bar{J} :

$$\bar{J} = \frac{\sum_{n=1}^N f^{(n)}(t_h^{(n)}) - u_m(d^{(n)}, t_m^{(n)})}{\sum_{n=1}^N t_h^{(n)} + t_m^{(n)}}. \quad [4]$$

The subject has control over two variables: t_m (duration of movement, how fast to move between patches) and t_h (duration of harvest, how long to stay at each patch). She should select these variables in such a way that \bar{J} is maximized. Typically, this is accomplished by finding the optimal duration $t_h^{(n)*}$ that the subject should stay at each patch. MVT (23) provides a solution for this objective. Here, our objective is broader: we aim to also find the optimum vigor of the movement via $t_m^{(n)*}$.

Whereas Charnov (23) assumed that the effort function u_m was linear in duration, here we assume that this function is concave upward. We think that this is justified, because the energetic expenditure during a movement is concave upward for many types of movements, such as walking (24), running (12), and reaching (13). This assumption is critical, because by doing so, we will be able to derive an expression for optimum vigor.

Suppose that the subject considers traveling to and harvesting in patch n . For this particular goal, we can rewrite Eq. 4 as follows:

$$\bar{J} = \frac{f^{(n)}(t_h^{(n)}) - u_m(d^{(n)}, t_m^{(n)}) + A}{t_m^{(n)} + t_h^{(n)} + B}. \quad [5]$$

In the above expression, A and B indicate terms in \bar{J} that do not depend on behavior associated with patch n . To predict behavior in patch n , we seek two unknowns $t_m^{(n)*}$ and $t_h^{(n)*}$, variables that maximize the above expression. We note that the derivative of \bar{J} with respect to the unknowns can be written in terms of \bar{J} :

$$\frac{d\bar{J}}{dt_m^{(n)}} = \left(\bar{J} + \frac{du_m}{dt_m^{(n)}} \right) \frac{-1}{t_m^{(n)} + t_h^{(n)} + B} \quad [6]$$

$$\frac{d\bar{J}}{dt_h^{(n)}} = \left(\frac{df^{(n)}}{dt_h^{(n)}} - \bar{J} \right) \frac{1}{t_m^{(n)} + t_h^{(n)} + B}.$$

The optimum vigor and harvest durations are found when these two equalities are simultaneously equal to zero. As a result, the optimum vigor is specified by a relationship between the effort expenditure during the movement and the global capture rate:

$$\frac{du_m}{dt_m^{(n)}} \Big|_{t_m^{(n)*}, t_h^{(n)*}} = -\bar{J} \Big|_{t_m^{(n)*}, t_h^{(n)*}}. \quad [7]$$

Furthermore, the optimum harvest duration is specified by a relationship between the harvest function and the global capture rate:

$$\frac{df^{(n)}}{dt_h^{(n)}} \Big|_{t_m^{(n)*}, t_h^{(n)*}} = \bar{J} \Big|_{t_m^{(n)*}, t_h^{(n)*}}. \quad [8]$$

To illustrate the meaning of these equations, suppose that one is harvesting reward at patch n . Harvesting commences when the subject arrives at the patch (Eq. 1), resulting in intake that rises (Fig. 1A). Harvesting requires effort, which according to Eq. 2, produces a peak in this function. If the objective was to maximize harvest intake at the current patch, the subject should stay until the time specified by the peak of this curve [i.e., when $df^{(n)}/dt_h^{(n)} = 0$]. However, Eq. 8 states that the harvest should end when the rate of intake equals the global capture rate in the environment, specified by \bar{J} . That is, the subject should abandon the patch when the rate of intake falls below the rate indicated by her past experience and future expectations. As a result, local factors (current reward, the effort required to harvest that reward) interact with global factors (past rewards, future efforts) to determine how long the subject should stay at the patch.

For example, the theory predicts that duration of stay at patch n should increase with reward available at that patch [effect of $\alpha^{(n)}$] (Fig. 1B) and decrease with effort required to acquire that reward [effect of $k^{(n)}$] (Fig. 1D). If past actions have been rewarding (increasing \bar{J}), the subject should exhibit impatience at the current patch, leaving the patch early (Fig. 1C). In contrast, if past actions have required large effort (decreasing \bar{J}), the subject should stay longer at the current patch (Fig. 1E). Finally, if future actions require high effort (decreasing \bar{J}) or are expected to encounter low reward, one should stay longer at the current patch.

More interestingly, the theorem predicts patterns of movement vigor. Movement requires energetic expenditure (function u_m). If we consider only the local conditions, then we would move in such a way as to minimize the energetic cost of travel [that is, find $t_m^{(n)}$ such that $du_m^{(n)}/dt_m^{(n)} = 0$]. This provides a ballpark estimate of how fast a movement takes place, as shown by those who study speed of locomotion in humans and other animals (11, 25, 26). However, our objective is to maximize the global utility \bar{J} . In that case, duration of movement will be modulated by the local contingencies and history of the subject. In particular, the duration of the movement toward patch n is optimal when the derivative of the energetic loss during the movement $du_m^{(n)}/dt_m^{(n)}$ is equal to $-\bar{J}$ (Fig. 1F). Therefore, the global capture rate (i.e., the history of the subject) affects both the duration of the harvest and vigor of the movement.

For example, the theory predicts that, when moving toward patch n , the subject will move faster if she expects greater reward

there (Fig. 1H) [increased future reward $\alpha^{(n)}$ increases \bar{J}] or expects to spend smaller effort there (Fig. 1I) [decreased future $k^{(n)}$ decreases \bar{J}]. Surprisingly, the theory also predicts that, when the subject is moving away from patch n , her movement will be faster if she has just completed a good harvest (Fig. 1H) [increased past reward $\alpha^{(n-1)}$ increases \bar{J}] or has just spent a small amount of effort harvesting that reward (Fig. 1I) [decreased $k^{(n-1)}$ decreases \bar{J}]. Therefore, the same factors that affect harvest duration, including past history and future expectations, should also affect movement vigor.

Alternatively, in contrast to Eq. 4, we can define the global utility as the average value of the individual utilities: $\bar{J} = N^{-1} \sum_{n=1}^N J^{(n)}$. In this case, the optimum harvest duration and movement speed depend only on the location conditions—history of reward and effort should not affect decision making and motor control. We set out to test these predictions.

Duration of Harvest Modulated Vigor of Movements. In foraging, reward is food, and its utility is calories acquired over the harvest period. In our foraging-like experiments, reward is information,

a finite amount of which is available in an image. As people gaze at an image, they are accumulating some amount of that reward. What is the utility of this act? That is, what is the harvest function $f(t_h)$ for this behavior? Because this question is central to the theory, we began with experiments that could, in principle, uncover the shape of $f(t_h)$.

If during gazing, the harvest function is concave downward, then presentation of two images in sequence during period $T = 2t_h$ will result in a greater harvest per unit time than just one image during the same period T (Fig. 2A). That is, a concave downward harvest function implies the following inequality:

$$f(\alpha, 2t_h) < 2f(\alpha, t_h). \quad [9]$$

This implies that, by controlling t_h , we are effectively modulating the rate of reward and therefore, modulating \bar{J} . In this scenario, reducing t_h should produce an increase in \bar{J} , predicting an increase in movement vigor.

However, if during gazing, the harvest function is linear with time, then changes in t_h should have no effect on movement vigor.

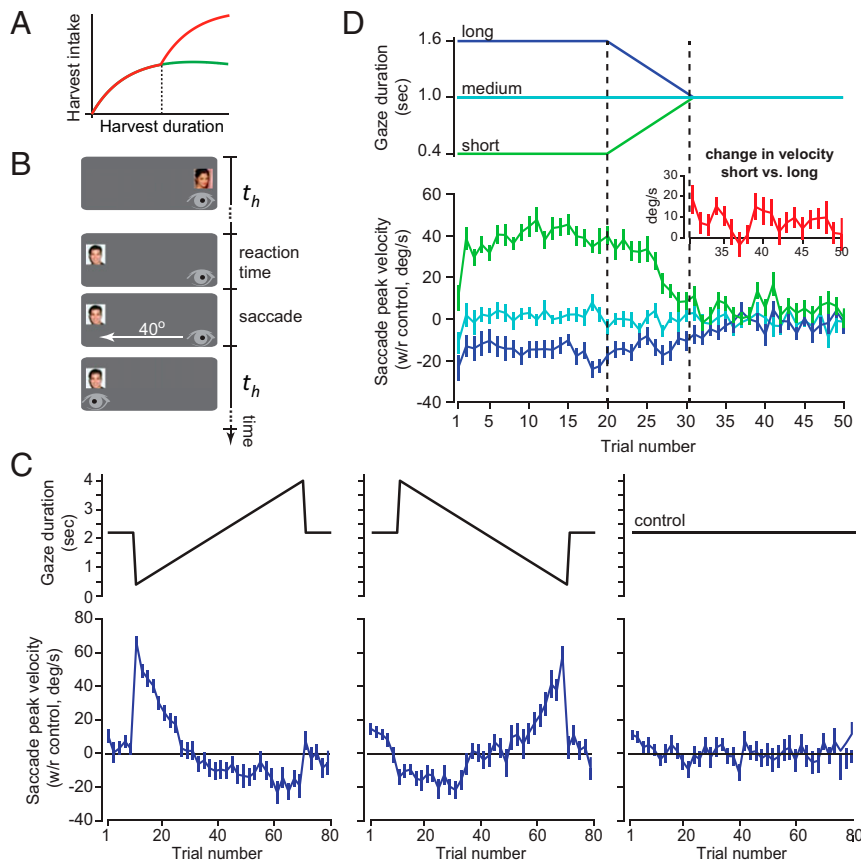


Fig. 2. Inferring the shape of the harvest function during image gazing. (A) If during image gazing, the harvest function is concave downward, then presentation of two images in sequence during period $T = 2t_h$ will result in a greater harvest per unit time than just one image during the same period T . As a result, reducing t_h should produce an increase in \bar{J} , predicting an increase in movement vigor. (B) Experiment design. We presented a small image of a face located at $\pm 20^\circ$ with respect to midline. After the subject made a saccade to it, we controlled duration t_h that they were allowed to gaze at that image before another image was presented. (C) Harvest duration strongly affected saccade vigor. We measured saccade peak velocity of each subject with respect to the average velocity during the control condition (constant t_h ; Right). An increase in t_h coincided with reduction in saccade peak velocity (Left). A reduction in t_h coincided with increase in saccade peak velocity (Center). Therefore, when time allowed for harvesting was short, people moved faster between the reward sites. (D) Harvest history modulated vigor. Subjects experienced a history of short, medium, or long harvest durations. They were then tested in identical harvest conditions (control trials 31–50). We measured within-subject change in saccade peak velocity with respect to the medium harvest trials. Saccade peak velocity was high during the short harvest trials and remained high in the control trials (Inset; within-subject change in peak velocity during control trials, short vs. long harvest). Therefore, experience of short harvests had long-lasting effects on saccade vigor. Data are mean \pm SEM.

Finally, if the harvest function is concave upward with time, then reductions in t_h should decrease movement vigor.

To test these alternatives, in Exp. 1 ($n = 16$ subjects), we manipulated t_h by controlling the amount of time that subjects were allowed to gaze at a small image of a face (Fig. 2B) located at $\pm 20^\circ$ with respect to midline. In some blocks of trials, t_h gradually increased, while in other blocks, t_h gradually decreased (or remained constant) (Fig. 2C). Decreasing t_h produced a robust increase in saccade velocities [repeated measures ANOVA, main effect of trial, $F(15,30) = 10.54, P < 2 \times 10^{-16}$], while increasing t_h produced a robust decrease in velocities [repeated measures ANOVA, main effect of trial, $F(15,30) = 8.52, P < 2 \times 10^{-16}$]. Similar changes were observed in saccade duration (SI Appendix, Fig. S1). That is, when the period allowed for gazing was short, the eyes moved faster between reward sites. These results revealed that, during image gazing, the harvest function $f(t_h)$ was concave downward.

Given this shape for the harvest function, theory predicted that, after a history of brief harvests t_h resulting in high global capture rate \bar{J} , vigor would remain high regardless of the local conditions (Fig. 1G). That is, vigor should be influenced by history of previous harvests. To test for this in Exp. 2 ($n = 19$ subjects), in some blocks, harvest duration t_h was long for 20 trials and then gradually decreased, while in other blocks, t_h was brief for 20 trials and then gradually increased. All blocks ended with control trials in which $t_h = 1$ s. Long harvests should produce a reduced rate of reward, resulting in smaller global capture rate \bar{J} . This should reduce vigor in subsequent control trials. We found that saccade velocity during the control trials was higher after a history of short harvests (Fig. 2D, Inset) [two-way repeated measures ANOVA, effect of group $F(1,36) = 9.382, P = 0.0041$].

Our results may have been influenced by factors unrelated to harvesting of reward. In particular, as we manipulated t_h , we also encouraged the subjects to increase the rate at which they made saccades. Was vigor simply related to rate of saccades, or was vigor associated with the amount of reward harvested during period t_h ? To answer this question, we manipulated the harvest function via reward α . Assuming that noise images are less valuable than faces, $\alpha_n < \alpha_f$, it follows that

$$f(\alpha_f, t_h) > f(\alpha_n, t_h). \quad [10]$$

Given some history of t_h , the harvested reward is always larger when α is larger, resulting in increased \bar{J} and predicting greater vigor despite an equal rate of saccades. To test for this, we had subjects (Exp. 1, $n = 16$) make saccades to either face or noise images (block design). In each block, we controlled t_h and found that, as expected, t_h modulated saccade velocity in both the face and the noise blocks. However, across values of t_h , velocity remained consistently higher in the face block compared with the noise block (SI Appendix, Fig. S1) [linear mixed effects model, increasing and decreasing blocks averaged together by harvest duration, main effect of image type $F(1,29) = 279.0, P < 10^{-3}$].

In summary, if during gazing, the harvest function $f(t_h)$ is concave downward, then two images in sequence produce a greater amount of harvest per unit of time than one image, leading to a larger global capture rate. We controlled t_h and observed that people responded by changing their saccade vigor: increasing vigor when gaze duration was short, and decreasing vigor when gaze duration was long. Furthermore, after experiencing a history of short gaze durations, subjects continued to move their eyes with greater vigor in control trials, illustrating that the past history of capture rates influenced current vigor.

Richness of the Environment Affected Harvest Duration and Saccade Vigor. In the natural setting, people control their harvest duration and movement vigor. An interesting prediction of the theory is that, in a rich environment where rewards are plentiful, subjects should reduce how much time they are willing to spend at any given site (Fig. 1C) (effect of \bar{J} on t_h) and then move with high vigor to the next site (Fig. 1H) (effect of \bar{J} on t_m).

In Exp. 3 ($n = 17$ subjects), the trials began with a center fixation that was followed by two simultaneous images, each selected randomly from five categories: (i) noise, (ii) simple shapes, (iii) inanimate objects, (iv) animate objects, and (v) faces (Fig. 3A). The images were always 20° apart, but their positions varied with respect to the midline. Subjects had 2 s to freely gaze (Fig. 3A). On each trial, we measured the time that they spent on each image and their peak saccade velocity as they moved their eyes from one image to another. Image category served to modulate reward magnitude α (with value increasing from noise to face), and image position served to modulate effort expenditure rate k (with k increasing with eccentricity). To analyze the data, we implemented a linear mixed effects model that related the dependent variables (gaze duration at the image, saccade peak velocity) to the independent variables (type of the image, eccentricity of the image, type of the other image).

To establish the relative value of each image, we quantified the probability of choosing that image after removal of the center fixation dot. This probability increased with image type (Fig. 3B) [$F(4,64) = 82.8, P < 10^{-15}$], suggesting that α increased from noise to face. To establish the relative effort associated with gazing at each image, we quantified choice probability as a function of image eccentricity. The probability decreased with image eccentricity (Fig. 3B). After the eyes landed on an image, gaze duration increased with image value (Fig. 3C, Left) [main effect of current image type, $F(4,400) = 56.3, P < 10^{-13}$]. That is, as reward magnitude at the current patch increased, so did harvest duration, confirming a prediction of the theory (Fig. 1B) [effect of $\alpha^{(n)}$ on $f^{(n)}$]. More interestingly, gaze duration decreased as the value of the competing image increased (Fig. 3C, Center) [main effect of competing image type, $F(4,400) = 46.8, P < 10^{-15}$]. That is, as the environment became richer, the amount of time devoted to harvesting any given image decreased (grays become lighter in each column from the top row to the bottom row in Fig. 3C, Right).

We expected that, as reward magnitude at the destination increased, saccade velocity toward that destination should also increase [effect of $\alpha^{(n+1)}$ on \bar{J}] (Fig. 1G). Indeed, saccade peak velocity increased with the value of the destination image (Fig. 3D, Left) [main effect of image type at destination, $F(4,400) = 7.42, P = 9 \times 10^{-6}$]. More interestingly, theory predicted that past history of reward should modulate movement vigor (Fig. 1H): after a rewarding event, \bar{J} increased, and therefore, people should move faster toward the next image. To test for this, we quantified the effect of the just viewed image on the saccade that took the gaze away from that image (i.e., toward the competing image) and found that, after viewing of a high-valued image, saccade peak velocity was greater as the eyes moved away (Fig. 3D, Center) [main effect of previously viewed image, $F(4,400) = 16.7, P = 1.1 \times 10^{-12}$]. As a result, saccade velocities were highest in the richest environment (Fig. 3D, Right).

Theory predicted that, as effort expenditure at the current patch increased, harvest duration should decrease [effect of $k^{(n)}$ on $f^{(n)}$] (Fig. 1D). We had assumed that image eccentricity would modulate effort expenditure associated with gazing. Indeed, time spent at the current image depended on the effort requirements of that image: as the image eccentricity increased, gaze duration decreased (Fig. 3E) [main effect of eccentricity, $F(8,128) = 6.91, P = 1.6 \times 10^{-7}$]. Furthermore, saccade velocities were highest when the two images were equally distant from center, declining as the asymmetry increased (Fig. 3F) [effect of difference in

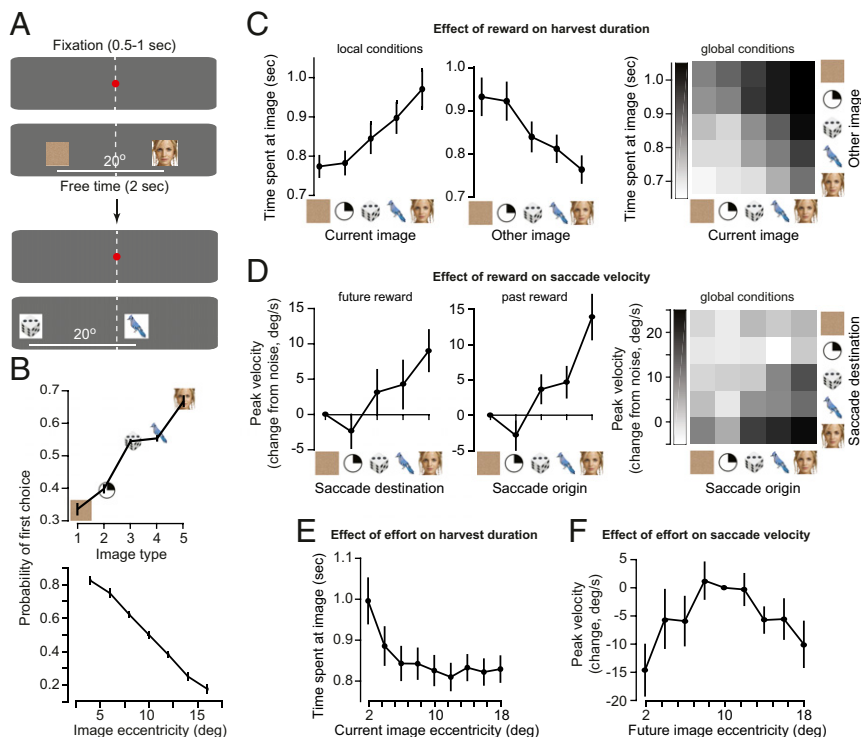


Fig. 3. Richness of the environment affected harvest duration and saccade vigor. (A) Trials began with presentation of two simultaneous images, each selected randomly from five categories. The two images were 20° apart, but their positions varied with respect to the midline. Subjects had 2 s to freely gaze. (B) Probability of choosing an image after removal of fixation dot. (C) Effect of image content (proxy for reward value) on harvest duration. Time spent gazing increased with current image value and decreased with competing image value. As the environment became richer (*Right*; from the top row to the bottom row of each column of the matrix), the amount of time spent at any given image decreased. (D) Effect of image content on saccade velocity. Saccade velocity was measured with respect to the noise condition. Velocity increased with image value at destination and also increased with the value of the image at the previous harvest. As the environment became richer (*Right*; from the top row to the bottom row of each column of the matrix), saccade peak velocity increased. (E) As image eccentricity increased, gaze duration tended to decrease. (F) Effect of eccentricity on saccade velocity. Saccade velocities were measured with respect to the condition where the two images were equally distant from center ($\pm 10^\circ$). Data are mean \pm SEM.

eccentricity of two images, $F(4,64) = 7.9$, $P = 3.1 \times 10^{-5}$. This last result is difficult to interpret, because as the future image becomes more eccentric (requiring more effort), the current image becomes less eccentric (requiring less effort). This experiment does not allow us to test the effects of effort expenditure during harvest on saccade vigor (Fig. 1*I*). We will explore this question directly in experiments below.

In summary, when subjects were presented with two reward patches, the duration of harvest at a given patch depended on both the magnitude of reward (image type) and the magnitude of effort required for harvesting that reward (eccentricity): people spent greater time gazing at more valuable images, but gaze duration decreased with increased effort requirements of harvest. Duration of harvest depended on global conditions: as the value of the competing image increased, people spent less time gazing at the current image. Reward available at the destination affected saccade vigor: increased image value at the destination produced increased vigor. However, past history of reward also affected vigor: saccade velocities were higher when subjects were leaving a high-valued image. As a result, richness of the environment affected both duration of harvest and vigor of movement. In a rich environment, people spent less time at the reward site and moved faster between the sites. All of these results were consistent with the theory.

History of Reward and Effort Affected Harvest Duration as Well as Saccade Vigor. In Exp. 3, saccades were made toward images, raising the possibility that low-level differences in the images (luminosity, contrast, etc.) might have affected saccade vigor.

Furthermore, that experiment imposed an explicit time limit on harvesting, wherein gazing at one image reduced the time available for gazing at the other image, thereby imposing an opportunity cost. For our next set of experiments, we designed a foraging-like environment in which there were no explicit time limits on harvesting, and all saccades were toward a single control stimulus.

In Exp. 4 ($n = 22$ subjects), subjects were presented with only one image at a time, allowing them to decide how long to gaze at that image without a time limit (Fig. 4*A*). We used three image types: (i) simple shapes, (ii) realistic objects, and (iii) faces. While subjects gazed at the image, they were presented with a dot that identified the location of a future image (located randomly at a distance of 10° , 15° , or 20°). After they made a saccade to the dot, they were presented with a randomly selected image at that location.

Theory predicted that harvest duration should increase with reward available at the current patch [effect of $\alpha^{(n)}$ on $f^{(n)}$] (Fig. 1*B*). To test for this, for each subject, we measured their average gaze duration \bar{t}_h across all images (633.8 ± 25 ms) and then represented gaze duration at the current image with respect to their \bar{t}_h . As expected, gaze duration increased with the amount of reward (Fig. 4*B*) [main effect of image type, $F(2,42) = 17.9$, $P = 2.3 \times 10^{-6}$]. This confirmed our results in Exp. 3, showing that, even without a competing image and time limits on harvest, increased current reward promoted harvest duration.

To test whether history of past rewards altered duration of current harvest, we computed gaze duration at current image as a function of the previously viewed image and found that, as the

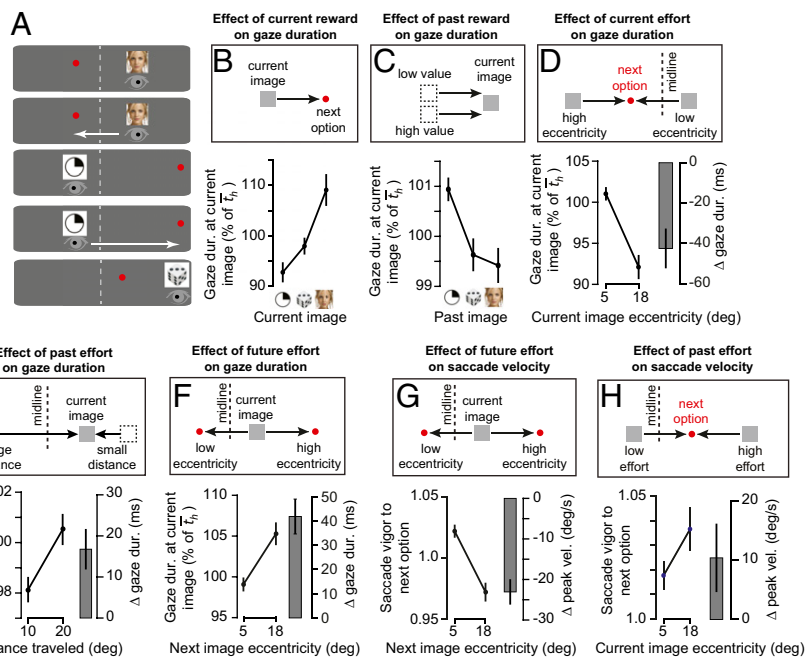


Fig. 4. Effect of reward and effort history on gaze duration. (A) Subjects were presented with a single image, allowing them to gaze without competition from another image and without a time limit. While they gazed, they were presented with a dot that identified the location of the future image. Only after making a saccade to the dot, they were presented with an image at that location. Images were selected randomly. (B) Gaze duration at current image normalized within subjects to their average gaze duration \bar{t}_h across all images. (C) Gaze duration at the current image was longer if the past image was less valuable. (D) Gaze duration at current image was shorter if that image had high eccentricity. *Lower Right* shows within-subject change in gaze duration. In this comparison, the next option was at the same eccentricity regardless of the current image eccentricity. (E) Gaze duration was longer if the effort required for travel to that image (distance) was larger. Note that previous image eccentricity was kept constant in this comparison. (F) Gaze duration at current image increased if future effort requirements (dot eccentricity) were high. (G) Saccade velocity was lower if the future harvest required large effort. Saccade vigor refers to amplitude-normalized peak velocity. (H) Saccades that followed a high-effort harvest were faster than saccades that followed a low-effort harvest. Data are mean \pm SEM. Horizontal error bars for all eccentricity figures are very small (0.2–0.4%).

value of the previous image increased, gaze duration at the current image decreased (Fig. 4C) [effect of image value at the previous trial on gaze duration in current trial, $F(2,42) = 9.97$, $P = 2.8 \times 10^{-4}$]. Therefore, after a poor harvest (decreased \bar{J}), people lingered longer at the next reward site.

Theory predicted that, if effort expenditure at the current patch was high, subjects should shorten their period of harvest [effect of $k^{(n)}$ on $f^{(n)}$] (Fig. 1D). We compared gaze duration at images of varying eccentricity while keeping the next option constant (Fig. 4D). As current image eccentricity increased, gaze duration decreased [effect of image eccentricity on gaze duration, $F(1,105) = 7.2$, $P = 1.6 \times 10^{-11}$].

Theory predicted that, if the previous harvest was effortful, reducing \bar{J} , then one should linger longer at the next harvest. Indeed, we found that, if the travel distance to the current image was large, greater time was spent at the current image (Fig. 4E) [effect of previous saccade distance on current gaze duration, $F(1,63) = 4.1$, $P = 0.046$]. If one expected to encounter high effort in the subsequent patch, one should increase harvest duration at the current patch. As the eccentricity of the future image increased, gaze duration at the current image increased (Fig. 4F) [effect of dot eccentricity on current gaze duration, $F(1,105) = 70.8$, $P = 2.1 \times 10^{-13}$]. Therefore, we found that past, current, and future effort expenditures all modulated harvest duration in the direction predicted by the theory.

We next asked whether effort requirements of harvest affected vigor of movements. Because in this experiment, the saccades were of various amplitudes and peak velocity varies with amplitude, we normalized the data for each saccade using procedures that took into account direction and amplitude of each movement (17, 27, 28). For each saccade, we computed vigor as the

ratio of its peak velocity with respect to the expected velocity for that subject given the amplitude and direction of the movement (*Methods*). The term vigor refers to this ratio.

Theory predicted that, when high effort was expected at a future location, one should make a low vigor movement toward that location [effect of $k^{(n+1)}$ on \bar{J}] (Fig. 1I). Indeed, saccade vigor toward the dot decreased as dot eccentricity increased (Fig. 4G) [effect of dot eccentricity on vigor, $F(1,105) = 16.5$, $P = 9.4 \times 10^{-5}$]. However, contrary to the theory's prediction and consistent with what we had observed in Exp. 3, after expenditure of high effort, the saccade to the next target was of high vigor (Fig. 4H) [effect of current image eccentricity on saccade vigor, $F(1,105) = 16.5$, $P = 9.4 \times 10^{-5}$].

In summary, in this experiment, subjects had unlimited time to stay and harvest reward. They were unaware of reward magnitude (image type) at the future harvest site but were provided with information regarding its effort requirements (eccentricity). We confirmed all of the theory's predictions regarding harvest duration: current reward increased harvest duration, while past reward decreased harvest duration. Current effort decreased harvest duration, while past and future effort increased harvest duration. Consistent with the theory, future expected effort decreased saccade vigor. However, once again, we found that, contrary to our expectations, past expenditure of high effort was followed by high vigor.

Past Effort Expenditure Increased Vigor. Why did our theory fail to predict the observed effects of past effort expenditure on vigor of movements? One possibility is that, as we expend effort harvesting reward, expenditure increases the subjective value of the reward. This phenomenon is termed “justification of effort,”

suggesting that subjective value of reward increases with the effort required to harvest it (29). If expenditure of effort during harvesting has increased the subjective value of reward, then given the opportunity to harvest again, subjects should move faster toward it and should stay longer there compared with when the opportunity comes after having spent low effort harvesting the same reward.

In our final experiment (Exp. 5, $n = 18$ subjects), we manipulated history of effort expenditure during harvest while repeatedly measuring vigor and harvest duration in probe trials. As before, subjects viewed an image while a dot provided them with information regarding the position of the next image. To modulate effort history, during a block of 100 trials, all images (except probe images) appeared at high eccentricity in the range from 10° to 30° with respect to midline (uniform distribution, mean eccentricity of 20°). In another block, the images appeared at low eccentricity within 10° of the midline (mean eccentricity of 0°). These blocks represented high-effort and low-effort environments, respectively. In each block, 25% of the trials were probe trials, placing the image exactly at 10° with respect to the midline and the dot also exactly at 10° with respect to midline. Images were selected randomly from two categories: faces and inanimate objects. As a result, the probe trials were identical in all blocks, allowing us to determine whether past effort expenditure modulated vigor and harvest duration.

We observed unequivocal effects: subjects spent greater time gazing at the probe images when those images appeared in the context of the high-effort environment (Fig. 5B) [effect of effort environment on gaze duration, two-way repeated measured ANOVA, $F(1,663) = 17.03$, $P = 4.15 \times 10^{-5}$]. Furthermore, saccade velocities in the probe trials from the image to the dot had higher velocity in the context of the high-effort environment (Fig. 5C) [effect of effort environment on saccade peak velocity, two-way repeated measured ANOVA, $F(1,663) = 37.78$, $P = 1.4 \times 10^{-9}$]. Therefore, if past harvests required high expenditure of effort, subjects behaved in probe trials as if their subjective value of reward had increased, lengthening their duration of gaze and then moving with greater vigor.

Discussion

During foraging, how long should one stay at a given site accumulating reward, and then how fast should one travel to the next site? Optimal foraging predicts that these behaviors aim to maximize a global utility: sum of all rewards acquired during the harvest minus all energy expended during travel and harvest periods divided by total time (i.e., the global capture rate). MVT provides a solution to the decision-making problem (23): the animal should leave the reward site when the local capture rate drops to the global capture rate. However, the theory assumes that “the length of time between patches should be independent of the length of time the predator hunts within any one.” That is, MVT assumes that factors that influence harvest duration do not affect movement vigor. Contrary to this assumption, experiments have found that increased reward at the destination shortens travel duration, increasing movement vigor (14, 15, 19, 30, 31). Because reward magnitude also modulates harvest duration, it seems likely that, during foraging, speed of travel and duration of harvest are not independent but influenced by the reward and effort characteristics of the environment. Here, we resolved this theoretical problem and tested the generalized theory in a foraging-like paradigm.

The generalized theory predicted that the subject should plan to arrive at her destination when the rate of energetic loss due to travel becomes equal to the negative of the average capture rate experienced in the environment. Furthermore, it predicted that variables that affected the global capture rate, like past rewards or future efforts, should influence both the harvest duration and movement vigor (Fig. 1 and *SI Appendix*, Table S1).

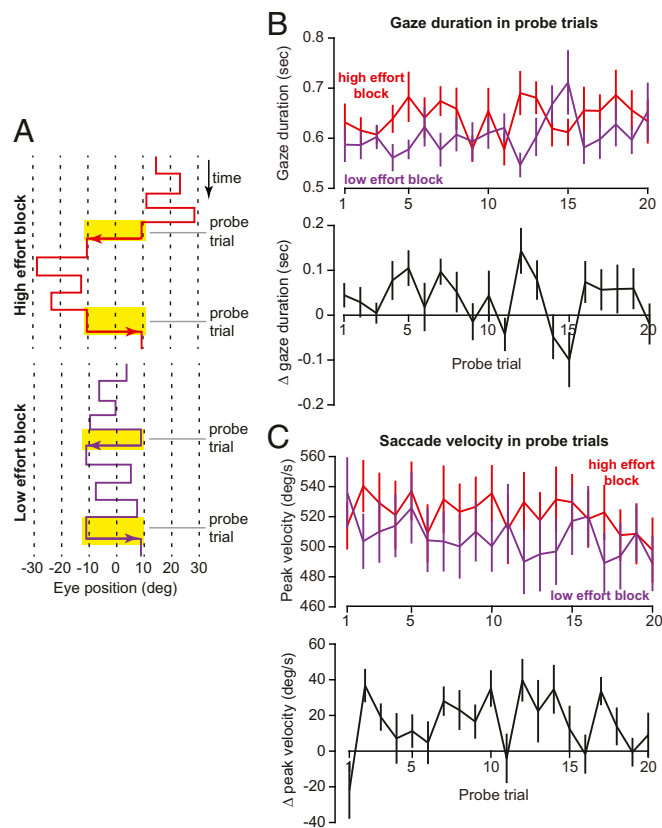


Fig. 5. History of effort affects harvest duration and saccade vigor. (A) Subjects were presented with an image at one location and a dot at another location. After making a saccade to the dot, they were presented with an image at that location. Effort history was manipulated via image eccentricity, and its effects were quantified in occasional probe trials in which image and dot locations were kept constant. The vertical traces show eye position during gazing, and the horizontal traces show saccades. In the high-effort block, the nonprobe images were located in the range from 10° to 30° . In the low-effort block, the nonprobe images were located within the $\pm 10^\circ$ range. (B) Gaze duration during probe trials in the low- and high-effort blocks and the across-block within-subject change in gaze duration. After expenditure of high effort, gaze duration in probe trials increased. (C) Probe trial saccade peak velocity in the low- and high-effort blocks and the across-block within-subject change in peak velocity. After expenditure of high effort, vigor in probe trials increased. Data are mean \pm SEM.

To test the predictions, we considered behavior in an environment where harvesting of reward was via gazing at images, reward magnitude was modulated via image content, and travel was via saccadic eye movements. Our rationale for assuming that viewing of images was akin to acquisition of reward was due to behavioral and neurophysiological responses that people have to images. People are willing to exert more effort (press keys) to view attractive female images compared with other images (32), and they make saccades that have higher velocities toward those images (31). Viewing of facial images activates some of the reward-related regions of the brain (33), with greater activation associated with greater facial attractiveness.

During gazing, effort expenditure was modulated via eccentricity of the image: holding the eyes requires activity in the oculomotor neurons, and the firing rates of these neurons during gazing increases with eccentricity (34–36). Eye eccentricity is encoded in discharge of foveal-related neurons in the frontal eye field (37). Given that a unit of energy is expended per action potential (38), the energy expenditure associated with holding the eyes is likely to increase with eccentricity.

As the theory had predicted, the duration of harvest at a given site (gaze duration) varied positively with the magnitude of reward (image type) and varied negatively with magnitude of effort required for harvesting that reward (eccentricity). Harvest duration also depended on the history of the subject: after a good harvest (high reward or low effort), subjects lingered for a shorter period of time in their next harvest. In anticipation of a bad harvest (high future effort), they lingered longer in their current image. These results confirmed all of the theory's predictions regarding decision making: current reward increased harvest duration, while past reward decreased harvest duration. Current effort decreased harvest duration, while past and future effort increased harvest duration.

Reward and effort also produced changes in vigor (saccade peak velocity). As reward available at the destination increased, so did vigor of the movement toward it. As effort required for a future harvest increased, vigor of the movement toward it decreased. More interestingly, after a good harvest, subjects increased their vigor as they left the harvest location. As a result, a history of high reward rates produced high vigor in control trials. All of these observations were in agreement with the theory.

However, contrary to our expectations, past expenditure of effort consistently produced increased saccade vigor. This may have been a reflection of elevation in the subjective value of reward after expenditure of effort, generally termed justification of effort (29, 39–43). That is, reward value α may depend on the history of past effort that the subject has expended in its acquisition. To test for this, we manipulated history of effort and then measured duration of harvest as well as vigor of movements during probe trials. We found that, if past harvests accompanied high-effort expenditure, in the following probe trials, subjects gazed for a longer period of time and moved with greater vigor, both of which are consistent with an increased valuation of reward. However, additional experiments are needed to explore this phenomenon. For example, one can divide images of similar content into classes that are associated with low or high history of past effort and then test whether that effort history modulates the subjective value of the specific class.

Previous work has approached the problem of vigor by proposing a utility in which movement duration discounts the value of reward, thereby encouraging vigor (13, 44–46). In these theories, vigor arises from a desire to maximize a local movement utility in which there is competition between the temporal cost of reward devaluation (move rapidly) and the velocity cost of effort expenditure (move slowly). However, because these theories describe a framework that is only concerned with the immediate movement, they have difficulty accounting for history-dependent changes in vigor. In contrast to these works, here we suggested that both the decision-making process regarding duration of harvest and the motor control process regarding vigor of movement may be performed via a comparison between a local capture rate, reflecting the reward and effort expenditure of the current action, and a global capture rate, reflecting the history and future expectations of the same variables.

Neural control of the decision regarding how long to stay and view an image partly depends on structures in the frontal lobe. During fixation, when animals decide between staying or making a saccade, foveal-related neurons in the frontal eye field (37) encode rate of effort expenditure associated with staying (i.e., eccentricity), while neurons in the cingulate cortex encode the value of leaving (7). As the animal continues fixation, cingulate activity rises to a threshold, at which point the decision is made to leave the patch, resulting in a saccade. The rate of rise is slower when the effort requirements of travel are larger (7) and faster when the environment has low rates of reward (47). Furthermore, the reward that is expected at the destination affects discharge of saccade-related neurons in the frontal eye field, rising faster and reaching a higher peak discharge in expectation

of greater reward (48). Together, the decision regarding harvest duration seems to be a comparison between a variable that reflects the local capture rate (reward available and effort required to harvest it) and a threshold. That threshold is affected by the reward history of the animal (47).

Neural control of vigor partly depends on structures in the basal ganglia. Substantia nigra pars reticulata (SNr), an output nucleus of the basal ganglia, inhibits the superior colliculus. More vigorous saccades are associated with a deeper pause in the firing rates of SNr cells (49), and reward modulates the depth of this pause (50). Within the basal ganglia, some cells in the caudate influence the discharge of SNr neurons directly, while other cells do so indirectly via their projections to the external segment of globus pallidus (GPe). Caudate cells receive dopamine projections and generally fire more before a rewarding saccade (51). Onset of a stimulus that promises reward results in a burst of dopamine (52), which is followed by a more vigorous saccade (53). Indeed, chronic reduction in the concentration of dopamine in the caudate reduces saccade vigor by around 30% (54). GPe cells inhibit SNr and fire more strongly preceding a more vigorous saccade, and bilateral lesion of this region eliminates the ability of the animal to modulate saccade vigor in response to changes in reward (53). Taken together, it seems that control of saccade vigor is partly associated with the amount of dopamine in the basal ganglia, modulating activity of caudate and affecting the depth of pause in the SNr.

Niv et al. (20) proposed that the reward history of the animal may be reflected in the tonic levels of dopamine. Pasquereau and Turner (55) found that, as monkeys became sated, there was a slow decline in the firing rate of dopamine neurons, coincident with a reduction in vigor of movements. However, other reports have not found a relationship between tonic discharge of dopamine cells and reward history (56, 57). Instead, reward history seems to be reflected in the tonic discharge of serotonergic neurons (57). More recent results show that, before a movement starts, the phasic discharge of certain dopamine neurons influences velocity of the upcoming movement (58). Therefore, while history of reward and effort seems to influence vigor, the neural basis of encoding that history remains poorly understood.

Our experiments had the advantage of producing thousands of decisions and movements in a short period of time. However, our approach had a number of disadvantages. We used images to represent reward patches and assumed that each image contained a finite amount of reward that was gradually acquired during gazing. A recent experiment in monkeys provides some evidence for this assumption (59). Unfortunately, no experiment has definitely established the shape of the harvest function during image gazing. Our theory needs to be further tested with new experiments in which the information content at each patch and its rate of acquisition are independently controlled.

We assumed that effort expenditure associated with a saccade was a concave upward function of duration. Firing rate of oculomotor neurons varies strongly with saccade velocity (35): as velocity decreases, the rate of action potentials decreases. However, we do not know whether an increase in saccade duration results in a reduction in the total number of action potentials produced by the muscles that move the eyes. This assumption also remains to be tested in experiments where cost of travel is objectively known.

Finally, foraging is typically studied in paradigms where travel is via locomotion or flight and harvesting is via accumulation of food, actions that take minutes to complete. In contrast, saccades and gazing consume fractions of a second. Does saving a few milliseconds of time matter? If we view gaze duration as a period in which the brain harvests information from the visual scene, the fact that we make two or three saccades each second during all waking hours makes each millisecond of travel and harvest relevant, consistent

with our observation that the saccadic system is highly sensitive to reward and effort parameters that affect foraging.

In summary, we presented a theory that could predict both the decision-making process regarding duration of harvest at the reward patch and the motor control process regarding speed of movement between the patches. The theory suggested that these behaviors relied on a comparison between the local capture rate and the global rate: the subject should move so that she arrives at the patch at the time when the rate of energetic loss becomes equal to the negative of the global capture rate. The subject should leave the patch when the rate of energetic gain becomes equal to the global capture rate. Over the course of five experiments, we observed that subjects varied their harvest duration and movement vigor in patterns that were largely in agreement with the theory. Taken together, it seems that both decision making and control of movement vigor during foraging may be accounted for by assuming that the brain maintains an account of two vari-

ables: the rate of gain associated with utility of the current action and the global rate of gain associated with previous actions.

Methods

Experimental procedures were approved by Johns Hopkins Institutional Review Board and all subjects signed an approved consent form. We measured motion of the eyes in $n = 92$ healthy humans (26.9 ± 8.6 y old, mean \pm SD) as they viewed images ($4 \times 4^\circ$) on a monitor. The measurements were via an infrared camera (sampling rate of 1 kHz). Across five experiments, we varied image content and location and measured how these variables affected duration of fixation and saccade vigor. To compute changes in saccade vigor, we used a maximum likelihood estimate of the within-subject peak velocity with respect to saccade amplitude (28). Saccade vigor was the ratio of the measured peak velocity with respect to the predicted velocity of that movement. Details are available in *SI Appendix*.

ACKNOWLEDGMENTS. Our work was supported by NIH Grants 1R01NS078311 and 1R01NS096083 and National Science Foundation Grant 1723967.

- Richardson H, Verbeek NAM (1986) Diet selection and optimization by northwestern crows feeding on Japanese littleneck clams. *Ecology* 67:1219–1226.
- Stephens DW, Krebs JR (1986) *Foraging Theory* (Princeton Univ Press, Princeton).
- Lemon WC (1991) Fitness consequences of foraging behaviour in the zebra finch. *Nature* 352:153–155.
- Bendesky A, Tsunozaki M, Rockman MV, Kruglyak L, Bargmann CI (2011) Catecholamine receptor polymorphisms affect decision-making in *C. elegans*. *Nature* 472:313–318.
- Cowie RJ (1977) Optimal foraging in great tits (*Parus major*). *Nature* 268:137–139.
- Bautista LM, Tinbergen J, Kacelnik A (2001) To walk or to fly? How birds choose among foraging modes. *Proc Natl Acad Sci USA* 98:1089–1094.
- Hayden BY, Pearson JM, Platt ML (2011) Neuronal basis of sequential foraging decisions in a patchy environment. *Nat Neurosci* 14:933–939.
- Nambodiri VM, Mihalas S, Marton TM, Hussain Shuler MG (2014) A general theory of intertemporal decision-making and the perception of time. *Front Behav Neurosci* 8:61.
- Constantino SM, Daw ND (2015) Learning the opportunity cost of time in a patch-foraging task. *Cogn Affect Behav Neurosci* 15:837–853.
- Diamond JS, Wolpert DM, Flanagan JR (2017) Rapid target foraging with reach or gaze: The hand looks further ahead than the eye. *PLoS Comput Biol* 13:e1005504.
- Ralston HJ (1958) Energy-speed relation and optimal speed during level walking. *Int Z Angew Physiol* 17:277–283.
- Stuedel-Numbers KL, Wall-Scheffler CM (2009) Optimal running speed and the evolution of hominin hunting strategies. *J Hum Evol* 56:355–360.
- Shadmehr R, Huang HJ, Ahmed AA (2016) A representation of effort in decision-making and motor control. *Curr Biol* 26:1929–1934.
- Opris I, Lebedev M, Nelson RJ (2011) Motor planning under unpredictable reward: Modulations of movement vigor and primate striatum activity. *Front Neurosci* 5:61.
- Manohar SG, et al. (2015) Reward pays the cost of noise reduction in motor and cognitive control. *Curr Biol* 25:1707–1716.
- Sackaloo K, Strouse E, Rice MS (2015) Degree of preference and its influence on motor control when reaching for most preferred, neutrally preferred, and least preferred candy. *OTJR (Thorofare, NJ)* 35:81–88.
- Reppert TR, Lempert KM, Glimcher PW, Shadmehr R (2015) Modulation of saccade vigor during value-based decision making. *J Neurosci* 35:15369–15378.
- Wang C, Xiao Y, Burdet E, Gordon J, Schweighofer N (2016) The duration of reaching movement is longer than predicted by minimum variance. *J Neurophysiol* 116:2342–2345.
- Summerside EM, Shadmehr R, Ahmed AA (2018) Vigor of reaching movements: Reward discounts the cost of effort. *J Neurophysiol* 119:2347–2357.
- Niv Y, Daw ND, Joel D, Dayan P (2007) Tonic dopamine: Opportunity costs and the control of response vigor. *Psychopharmacology (Berl)* 191:507–520.
- Guitart-Masip M, Beierholm UR, Dolan R, Duzel E, Dayan P (2011) Vigor in the face of fluctuating rates of reward: An experimental examination. *J Cogn Neurosci* 23:3933–3938.
- Wilson RS, Husak JF, Halsey LG, Clemente CJ (2015) Predicting the movement speeds of animals in natural environments. *Integr Comp Biol* 55:1125–1141.
- Charnov EL (1976) Optimal foraging, the marginal value theorem. *Theor Popul Biol* 9:129–136.
- Zarrugh MY, Todd FN, Ralston HJ (1974) Optimization of energy expenditure during level walking. *Eur J Appl Physiol Occup Physiol* 33:293–306.
- Hoyt DF, Taylor CR (1981) Gait and the energetics of locomotion in horses. *Nature* 292:239–240.
- Selinger JC, O'Connor SM, Wong JD, Donelan JM (2015) Humans can continuously optimize energetic cost during walking. *Curr Biol* 25:2452–2456.
- Choi JE, Vaswani PA, Shadmehr R (2014) Vigor of movements and the cost of time in decision making. *J Neurosci* 34:1212–1223.
- Reppert TR, et al. (2018) Movement vigor as a traitlike attribute of individuality. *J Neurophysiol* 120:741–757.
- Clement TS, Feltus JR, Kaiser DH, Zentall TR (2000) "Work ethic" in pigeons: Reward value is directly related to the effort or time required to obtain the reward. *Psychon Bull Rev* 7:100–106.
- Takikawa Y, Kawagoe R, Itoh H, Nakahara H, Hikosaka O (2002) Modulation of saccadic eye movements by predicted reward outcome. *Exp Brain Res* 142:284–291.
- Xu-Wilson M, Zee DS, Shadmehr R (2009) The intrinsic value of visual information affects saccade velocities. *Exp Brain Res* 196:475–481.
- Aharon I, et al. (2001) Beautiful faces have variable reward value: fMRI and behavioral evidence. *Neuron* 32:537–551.
- O'Doherty J, et al. (2003) Beauty in a smile: The role of medial orbitofrontal cortex in facial attractiveness. *Neuropsychologia* 41:147–155.
- Robinson DA (1970) Oculomotor unit behavior in the monkey. *J Neurophysiol* 33:393–403.
- Sylvestre PA, Cullen KE (1999) Quantitative analysis of abducens neuron discharge dynamics during saccadic and slow eye movements. *J Neurophysiol* 82:2612–2632.
- Shadmehr R (2017) Distinct neural circuits for control of movement vs. holding still. *J Neurophysiol* 117:1431–1460.
- Segraves MA (1992) Activity of monkey frontal eye field neurons projecting to oculomotor regions of the pons. *J Neurophysiol* 68:1967–1985.
- Attwell D, Laughlin SB (2001) An energy budget for signaling in the grey matter of the brain. *J Cereb Blood Flow Metab* 21:1133–1145.
- Aronson E, Mills J (1959) The effect of severity of initiation on liking for a group. *J Abnorm Soc Psychol* 59:177–181.
- Neuringer AJ (1970) Many responses per food reward with free food present. *Science* 169:503–504.
- Kacelnik A, Marsh B (2002) Cost can increase preference in starlings. *Anim Behav* 63:245–250.
- Klein ED, Bhatt RS, Zentall TR (2005) Contrast and the justification of effort. *Psychon Bull Rev* 12:335–339.
- Zentall TR (2013) Animals prefer reinforcement that follows greater effort: Justification of effort or within-trial contrast? *Comp Cogn Behav Rev* 8:60–77.
- Shadmehr R, Orban de Xivry JJ, Xu-Wilson M, Shih TY (2010) Temporal discounting of reward and the cost of time in motor control. *J Neurosci* 30:10507–10516.
- Rigoux L, Guigon E (2012) A model of reward- and effort-based optimal decision making and motor control. *PLoS Comput Biol* 8:e1002716.
- Berret B, Jean F (2016) Why don't we move slower? The value of time in the neural control of action. *J Neurosci* 36:1056–1070.
- Barack DL, Chang SWC, Platt ML (2017) Posterior cingulate neurons dynamically signal decisions to disengage during foraging. *Neuron* 96:339–347.e5.
- Glaser JJ, et al. (2016) Role of expected reward in frontal eye field during natural scene search. *J Neurophysiol* 116:645–657.
- Sato M, Hikosaka O (2002) Role of primate substantia nigra pars reticulata in reward-oriented saccadic eye movement. *J Neurosci* 22:2363–2373.
- Handel A, Glimcher PW (1999) Quantitative analysis of substantia nigra pars reticulata activity during a visually guided saccade task. *J Neurophysiol* 82:3458–3475.
- Kawagoe R, Takikawa Y, Hikosaka O (1998) Expectation of reward modulates cognitive signals in the basal ganglia. *Nat Neurosci* 1:411–416.
- Matsumoto M, Hikosaka O (2007) Lateral habenula as a source of negative reward signals in dopamine neurons. *Nature* 447:1111–1115.
- Tachibana Y, Hikosaka O (2012) The primate ventral pallidum encodes expected reward value and regulates motor action. *Neuron* 76:826–837.
- Kori A, et al. (1995) Eye movements in monkeys with local dopamine depletion in the caudate nucleus. II. Deficits in voluntary saccades. *J Neurosci* 15:928–941.
- Pasquereau B, Turner RS (2013) Limited encoding of effort by dopamine neurons in a cost-benefit trade-off task. *J Neurosci* 33:8288–8300.
- Matsumoto M, Hikosaka O (2009) Two types of dopamine neuron distinctly convey positive and negative motivational signals. *Nature* 459:837–841.
- Cohen JY, Amoroso MW, Uchida N (2015) Serotonergic neurons signal reward and punishment on multiple timescales. *eLife* 4:e06346.
- da Silva JA, Tecuapetla F, Paixão V, Costa RM (2018) Dopamine neuron activity before action initiation gates and invigorates future movements. *Nature* 554:244–248.
- Turrin C, Fagan NA, Dal Monte O, Chang SWC (2017) Social resource foraging is guided by the principles of the marginal value theorem. *Sci Rep* 7:11274.