

# Neuronal basis of sequential foraging decisions in a patchy environment

Benjamin Y Hayden<sup>1,2</sup>, John M Pearson<sup>1,2</sup> & Michael L Platt<sup>1-3</sup>

Deciding when to leave a depleting resource to exploit another is a fundamental problem for all decision makers. The neuronal mechanisms mediating patch-leaving decisions remain unknown. We found that neurons in primate (*Macaca mulatta*) dorsal anterior cingulate cortex, an area that is linked to reward monitoring and executive control, encode a decision variable signaling the relative value of leaving a depleting resource for a new one. Neurons fired during each sequential decision to stay in a patch and, for each travel time, these responses reached a fixed threshold for patch-leaving. Longer travel times reduced the gain of neural responses for choosing to stay in a patch and increased the firing rate threshold mandating patch-leaving. These modulations more closely matched behavioral decisions than any single task variable. These findings portend an understanding of the neural basis of foraging decisions and endorse the unification of theoretical and experimental work in ecology and neuroscience.

Resources are rarely distributed uniformly in the environment. Food, water and other vital commodities more often occur in spatially localized and temporally ephemeral patches<sup>1</sup>. Patchy environments force animals to balance the benefits of staying in a depleting patch and leaving for a richer one<sup>2</sup>. According to the marginal value theorem (MVT) of behavioral ecology, animals should leave patches when their intake rate diminishes to the average intake rate for the overall environment<sup>2,3</sup>. Organisms as diverse as worms, bees, wasps, spiders, fish, birds, seals and even plants obey the MVT<sup>3-6</sup>. Ethnographic evidence demonstrates that human subsistence foragers also obey the predictions of the MVT in their hunting behavior<sup>7</sup>, and laboratory findings suggest that monkeys may do so as well<sup>8</sup>. The generality of the MVT solution to the patch-leaving problem suggests that the underlying mechanism is fundamental to the way organisms make decisions<sup>4</sup>. The neuronal basis of patch-leaving decisions, however, remains unknown.

Building on recent progress toward understanding the neuronal mechanisms mediating perceptual decisions<sup>9</sup>, we hypothesized that the brain maintains a decision variable specifying the current relative value of leaving a patch. Conceptually, a decision variable is an analog quantity that incorporates all sources of information—in this case, reward size, handling time, search time and travel time—evaluated by the decision policy to generate a behavioral choice<sup>9</sup>. The hypothesized decision variable gives rise to a decision via comparison with a specific threshold. For simplicity, we assume that this process is analogous, although not necessarily isomorphic, to the neural integrate-to-threshold processes thought to mediate perceptual judgments<sup>9-13</sup>. We further conjecture that travel time between patches influences leaving decisions by changing the rate at which the decision variable grows, the threshold or both<sup>10,11</sup>.

To test these hypotheses, we developed a virtual foraging task in which rhesus monkeys chose one of two targets. One target

corresponded to remaining in the patch, and choosing it yielded a juice reward that declined each time it was chosen. The other target corresponded to leaving the patch, and choosing it yielded only a delay before the opportunity to choose again at a replenished patch. Monkeys' behavior closely matched the predictions of the MVT. We then recorded activity of neurons in the dorsal anterior cingulate cortex (dACC) while they performed the task.

dACC has been linked to reward outcome monitoring and behavioral adjustment<sup>14-16</sup>, as well as to signaling reward outcomes and predicting changes in behavior<sup>17-24</sup>. Notably, ACC dysfunction attends clinical disorders that are associated with difficulty in abandoning maladaptive patterns of behavior or cognition, including depression, addiction, obsessive-compulsive disorder and Tourette Syndrome<sup>25-27</sup>.

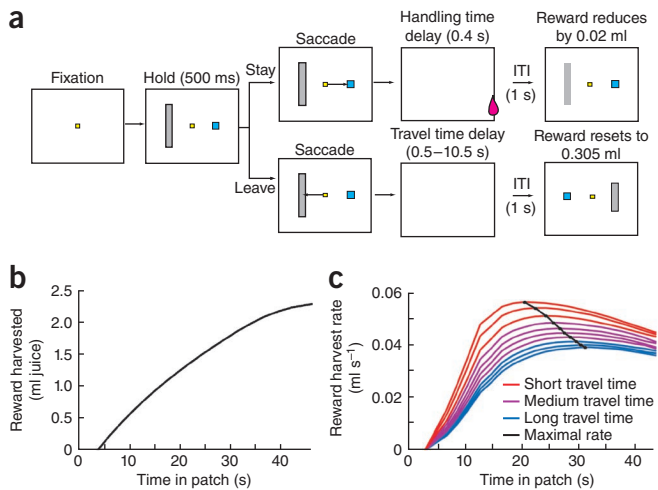
We found that dACC neurons responded each time monkeys made a choice and that these responses increased with time spent in the current patch. Monkeys abandoned a patch when neuronal responses reached a threshold associated with a particular travel time. When travel time between patches was high, the gain of neuronal responses with each decision to remain in the patch was smaller and the threshold for patch abandonment was higher than when travel time was short. Overall, neuronal response gain and threshold jointly predicted patch-leaving decisions. These findings suggest that dACC mediates patch-leaving decisions using a common integrate-to-threshold mechanism.

## RESULTS

For each choice there were two options (Fig. 1a). Choosing the stay (short blue) target led to a juice reward in 0.4 s (handling time). The value of this reward declined by  $19 \mu\text{l} \pm \epsilon$  (s.e.m.,  $\epsilon = 1.9 \mu\text{l}$ ) each time it was chosen, mirroring the diminishing returns common to patchy foraging environments (Fig. 1b). Choosing the leave (tall gray) target led to no reward and a long delay that was fixed in a patch and varied

<sup>1</sup>Department of Neurobiology, Duke University School of Medicine, Durham, North Carolina, USA. <sup>2</sup>Center for Cognitive Neuroscience, Duke University, Durham, North Carolina, USA. <sup>3</sup>Department of Evolutionary Anthropology, Duke University, Durham, North Carolina, USA. Correspondence should be addressed to B.Y.H. (benhayden@gmail.com).

Received 9 November 2010; accepted 28 March 2011; published online 5 June 2011; doi:10.1038/nn.2856



**Figure 1** Patch-leaving task. **(a)** Task design. After fixation, two eccentric targets, a large gray and a small blue rectangle, appear. Monkey chooses one of two targets by shifting gaze to it. Choice of blue rectangle (stay in patch) yields a short delay (0.4 s, handling time) and reward whose value diminishes by 19  $\mu$ l per trial. Choice of gray rectangle (leave patch) yields no reward and a long delay (travel time) whose duration is indicated by the height of the bar, and resets the value of the blue rectangle at 306  $\mu$ l. Travel time varies randomly from patch to patch and ranges from 0.5 to 10.5 s. **(b)** Plot of the cumulative reward available in this task as a function of time in patch, given the search times associated with animals' performance in the task (black line). Data are generated on the basis of average times associated with performance. **(c)** Plot of reward intake rate derived from a range of patch residence times (x axis: range of residence times). Data are shown for each of ten travel times (1-s intervals from 0.5 to 10.5 s). Rate-maximizing time in patch (the curves' maxima, shown by the black line) increases with increasing travel time. Data are generated based on average times associated with actual animal performance.

between patches, and reset the value of the stay target to its initial high value (306  $\mu$ l). We defined search time as any additional time spent in the patch not explicitly waiting and residence time as total time from arrival at a new patch, including handling times, search time and intertrial intervals (ITIs). Travel time was explicitly cued on all trials by the height of the gray bar and was reset to a new random value each time it was chosen (0.5 to 10.5 s, uniform distribution). The blue and gray bars alternated sides each time the leave target was chosen; any potential laterality in neural tuning functions was assumed to average out (**Supplementary Data 1**). In contrast to some natural foraging decisions, there was no physical travel during the travel time, nor was any action required during this delay; the only explicit cost of delay was opportunity cost. These simplifications are, we believe, not critical, as most other laboratory foraging tasks eschew effort requirements (for example, see refs. 8,28).


### Monkeys approximate rate maximization according to MVT

As travel time between patches grows, so does the rate-maximizing residence time (**Fig. 1c**). Consistent with the MVT, monkeys' patch-residence times rose with increasing travel time and were nearly rate maximizing ( $P < 0.0001$ ,  $\beta = 1.247$ , regression of residence time (s) against travel time (s); **Fig. 2a**). These effects were found in both monkeys individually ( $P < 0.0001$  for both individuals,  $\beta = 1.11$  for monkey E,  $\beta = 1.47$  for monkey O; **Supplementary Data 2** and **Supplementary Fig. 1**). Overall, both monkeys obtained 97.2% of the reward obtained by the best-fit rate-maximization algorithm (note that this is a measure of reward obtained versus maximal obtainable, not a measure of variance in behavior explained). Both monkeys remained

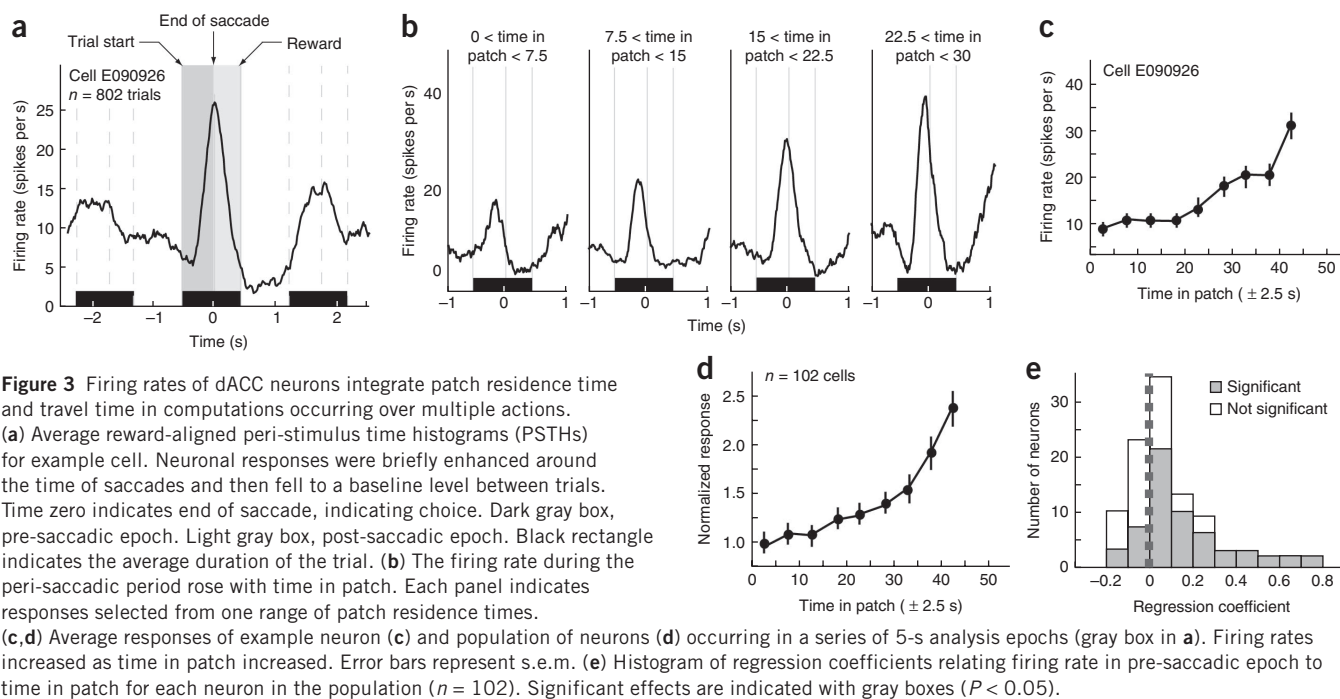
in patches slightly longer than predicted by the MVT (mean 2.2 s longer,  $P < 0.01$ , Student's  $t$  test). This slight over-staying may reflect a weak preference for immediate small rewards over delayed large rewards<sup>29–31</sup>, a slight over-estimate of travel times or even a status quo bias<sup>32</sup>. Leaving time was not influenced by travel time on the previous patch (regression of residence time against previous travel time,  $P = 0.44$ ; **Supplementary Data 3** and **Supplementary Fig. 2**).

Monkeys attempting to maximize local intake rates over the long-term should consider handling time as well as travel time<sup>2,3</sup>. To confirm that monkeys do so, we performed an additional behavioral experiment in which handling times, but not travel times, were varied from patch to patch (11 sessions, 6 in monkey E, 5 in monkey O). In each patch, handling time took one of ten values: 0.1, 0.4, 0.6, 1.1, 1.6, 2.1, 2.6, 3.1, 3.6 or 4.1 s. We cued handling time by varying the height of the blue rectangle; travel time was held constant at 5 s. We performed these experiments after the monkeys had learned the task, but before we began physiological data collection; monkeys received one full day of training each. As predicted, patch residence times declined with increased handling time (regression,  $\beta = -3.71$  for both,  $-3.81$  for E and  $-3.62$  for O,  $P < 0.001$  in all cases; **Fig. 2b**). Leaving times did not differ systematically from the rate-maximizing predictions at any of the ten points ( $P > 0.05$  in all cases, Student's  $t$  test). As an additional control, during these sessions, we interleaved standard fixed handling time patches with variable travel times. The average residence time on these trials was consistent with those obtained in the handling time control (**Fig. 2b**).

A natural question is whether the monkeys' foraging behavior may be explained in a delay-discounting framework<sup>30,33</sup>. In such a framework, each leave/stay decision is regarded as a choice between a smaller-sooner stay reward and larger-later leave reward (that is, the first reward in the new patch). Such a decision model would naturally account for the monkeys' observed tendency to stay longer in patches when faced with longer travel times, as the delay associated with patch-leaving would lead to discounting of the larger-later reward. To test this idea, we compared an empirically derived sequential foraging model inspired by the MVT against a standard delay-discounting model in which the hyperbolic discount parameter  $k$



**Figure 2** Monkeys obey the marginal value theorem in a virtual patchy foraging task. **(a)** Monkeys remain in the patch longer as travel time rises, as predicted by the marginal value theorem (MVT). Each dot indicates a single patch-leaving decision ( $n = 2,834$  patch-leaving events). The time at which the monkey chose to leave the patch (y axis) was defined relative to the beginning of foraging in that patch. Travel time was kept constant in a patch (x axis). Data from both monkeys is shown. Behavior (average is traced by the blue line) closely followed the rate-maximizing leaving time (red line), albeit delayed by 0–2 s. **(b)** Performance of two monkeys on handling time variant of patch-leaving task. In this control experiment, travel time was held constant (5 s) and handling time was randomly reset between each patch to have one of ten values. Patch residence time fell as handling time rose, consistent with the MVT. Observed times are shown with black dots; averages are shown with solid blue line. Best-fit line (dashed blue line) is nearly identical to rate-maximizing (red line). Average patch residence time on the interleaved standard travel time version of the task was consistent with this curve as well (red dots).



**Figure 3** Firing rates of dACC neurons integrate patch residence time and travel time in computations occurring over multiple actions. **(a)** Average reward-aligned peri-stimulus time histograms (PSTHs) for example cell. Neuronal responses were briefly enhanced around the time of saccades and then fell to a baseline level between trials. Time zero indicates end of saccade, indicating choice. Dark gray box, pre-saccadic epoch. Light gray box, post-saccadic epoch. Black rectangle indicates the average duration of the trial. **(b)** The firing rate during the peri-saccadic period rose with time in patch. Each panel indicates responses selected from one range of patch residence times. **(c,d)** Average responses of example neuron **(c)** and population of neurons **(d)** occurring in a series of 5-s analysis epochs (gray box in **a**). Firing rates increased as time in patch increased. Error bars represent s.e.m. **(e)** Histogram of regression coefficients relating firing rate in pre-saccadic epoch to time in patch for each neuron in the population ( $n = 102$ ). Significant effects are indicated with gray boxes ( $P < 0.05$ ).

was estimated by maximum likelihood (best fit,  $k = 1.26 \text{ s}^{-1}$ ). The Akaike weights for the two models (a measure of goodness of fit that accounts for different numbers of parameters in different models;  $w_{\text{MVT}} =$  weight for MVT model,  $w_{\text{DD}} =$  weight for delay-discounting model) clearly favored the sequential-trial foraging model, endorsing the MVT account of decision making in our task ( $w_{\text{MVT}}/w_{\text{DD}} = 1.31 \times 10^{52}$ ) (Supplementary Data 4 and Supplementary Fig. 3).

### Neurons integrate information over multiple actions

We recorded the activity of 102 single neurons in dACC in two monkeys performing this task (52 neurons in monkey E, 50 in monkey O; Fig. 3). For an example neuron, neural activity was aligned to the end of the choice saccade (time zero; Fig. 3a). Firing rate rose to a peak around the time of the choice saccade and then returned to a baseline value between trials. Such brief, peri-saccadic responses, often modulated by reward size and task context, are characteristic of neurons in dACC<sup>18,21,22,24</sup>. We focused on neuronal activity in the 500-ms epoch preceding saccade onset (pre-saccadic epoch). For most analyses, we focused on neural data associated with choosing to remain in the patch and excluded neural data associated with choosing to leave the patch (exceptions are noted). Data for individual subjects matched the combined data (Supplementary Data 5 and Supplementary Fig. 4).

We next examined the responses of the example neuron from four time periods relative to the beginning of foraging in the patch ( $t < 7.5 \text{ s}$ ,  $7.5 < t < 15$ ,  $15 < t < 22.5$ , and  $22.5 < t$ ; Fig. 3b). For this neuron, responses rose with cumulative time spent foraging in the patch. To quantify this enhancement, we measured pre-saccadic responses in a series of non-overlapping 5-s time bins (Fig. 3c). We included in each time bin all of the decisions in which the end of the saccade occurred in that bin. We found that firing rates rose with increasing patch residence time ( $\beta = 0.31$ ,  $P < 0.0001$ , linear regression of firing rate (spikes per s) against time in patch (s)). The same effects were observed in the population average firing rates ( $\beta = 0.18$ ,  $P < 0.0001$ , regression; Fig. 3d, Supplementary Data 6 and Supplementary Fig. 5). We observed a significant ( $P < 0.05$ ) positive regression coefficient in 49 neurons (average  $\beta = 0.24$  in significantly modulated cells),

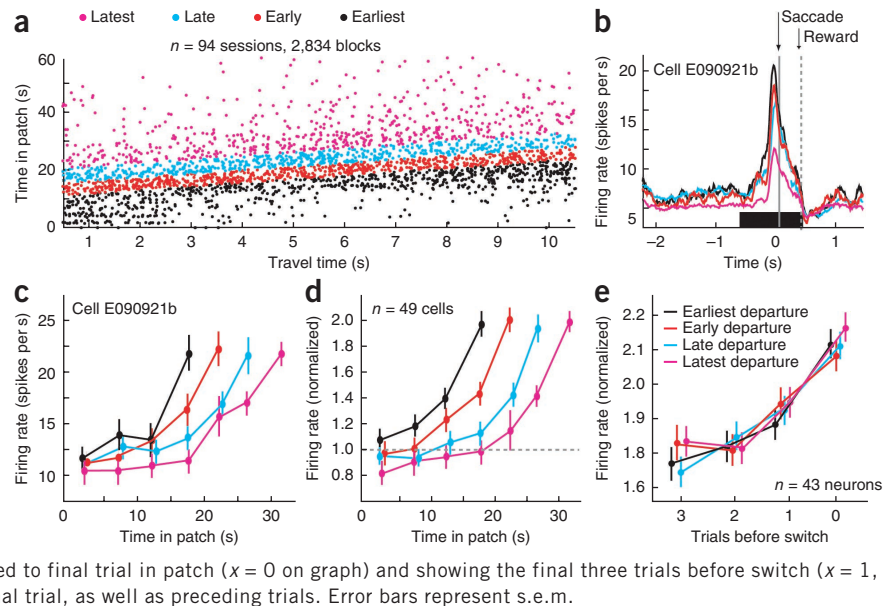
a significant negative slope in 10 (average  $\beta = -0.09$ ) and no significant slope in the remainder ( $P > 0.05$ ,  $n = 43$ , average  $\beta = 0.041$ ). The 49 neurons with positive slopes constitute the focus of subsequent analyses (Supplementary Data 7 and Supplementary Figs. 6 and 7).

We next performed the same analysis on two later epochs. In the post-saccadic epoch, we measured firing rates during the 400-ms handling time period beginning at saccade termination and ending with the reward. In the ITI epoch, we measured firing rates during the 1-s period beginning after reward delivery and ending when the next set of choice options was presented to the monkey. For the post-saccadic epoch, we observed a positive regression coefficient in 44 neurons ( $P < 0.05$ , average  $\beta = 0.15$  in significantly modulated cells), a negative coefficient in 8 (average  $\beta = -0.1$ ), and no significant effect in the remainder ( $P > 0.05$ , 50 neurons, average  $\beta = 0.010$ ). Nor was there much evidence of an effect in the ITI epoch at the population level; we observed a significant correlation between firing rates in the ITI and patch residence time in only four neurons (average  $\beta = 0.02$ ), all positive in sign. These numbers are similar to what would be expected by chance (Supplementary Data 8). This result suggests that dACC does not maintain a representation of time spent foraging in a patch across multiple actions; the locus of this trace in the brain remains to be determined. Overall, we observed weak or no effect of saccade direction on responses (Supplementary Data 1).

### Threshold-crossing of dACC firing predicts patch-leaving

The gradual rise in neural responses across decisions to stay in a patch resembles the within-trial rise-to-threshold processes observed in lateral intraparietal area, frontal eye fields (FEFs) and superior colliculus during motor preparation and decision-making<sup>10–12,34,35</sup>. We wondered whether a similar rise-to-threshold model might also account for the relationship between firing rates in dACC and patch-leaving decisions. To test this idea, we performed an analysis modeled on a previously developed method<sup>11</sup> for probing the relationship between the firing rates of FEF neurons and saccade initiation. Although FEF firing rates in that study<sup>11</sup> rose gradually to a fixed threshold on a single trial, the analogous rise in our study

**Figure 4** Firing rates of dACC neurons rise to a threshold associated with patch abandonment. **(a)** Plot of patch-leaving times, separated by whether they were earlier or later than the average leaving time. We divided patch-leaving decisions into four categories: earliest (black), early (red), late (cyan) and latest (magenta). These variables are independent of travel time and time in patch, meaning that, for example, earliest trials are equally likely to occur at any travel time (x axis) and any time in patch (y axis). **(b)** PSTH for an example neuron separated by earliness level. dACC neurons responded sooner and more strongly on earlier trials than on later trials. Black rectangle indicates the average duration of the trial. **(c,d)** Average firing rates of example neuron **(c)** and population **(d)** separated by earliness level. Firing rates rose faster for earlier patches but asymptoted at the same level. Error bars represent s.e.m. **(e)** Plot of average firing rate of population of neurons, aligned to final trial in patch ( $x = 0$  on graph) and showing the final three trials before switch ( $x = 1, 2$  and 3). Firing rates rose to the same level on final trial, as well as preceding trials. Error bars represent s.e.m.



was associated with changes in firing rate occurring in discrete bouts over multiple actions. Moreover, firing rates of FEF neurons gradually rise to threshold over tens of milliseconds, whereas the amplitudes of discrete dACC neuronal responses in our task increased over tens of seconds, orders of magnitude longer. Nonetheless, these analytical methods in principle generalize readily to our task. Our analysis asked two questions. First, does variability in patch-leaving times correlate with variability in the rate at which neural activity rises? Second, do firing rates rise to the same level regardless of the precise time monkeys choose to leave the patch for a given travel time?

We first divided all patch-leaving choices into residence time quartiles in each travel time (medians of each set for the aggregate response, earliest = 14.1 s, early = 19.2 s, late = 23.5 s, latest = 32.2 s; **Fig. 4a**). We refer to this classification of trials as the ‘earliness’ for each patch. We repeated the classification of leaving times into four different earliness bins for each neuron separately. By design, earliness is orthogonal to travel time, and so neural correlates of earliness and travel time are independent. For our example neuron, firing rates were significantly higher for the earliest than for the early patches (difference = 1.8 spikes per s,  $P = 0.02$ ; **Fig. 4b**), significantly higher for the early than for the late patches (difference = 2.0 spikes per s,  $P < 0.01$ ) and significantly higher for the late than for the latest patches (difference = 4.7 spikes per s,  $P < 0.01$ ).

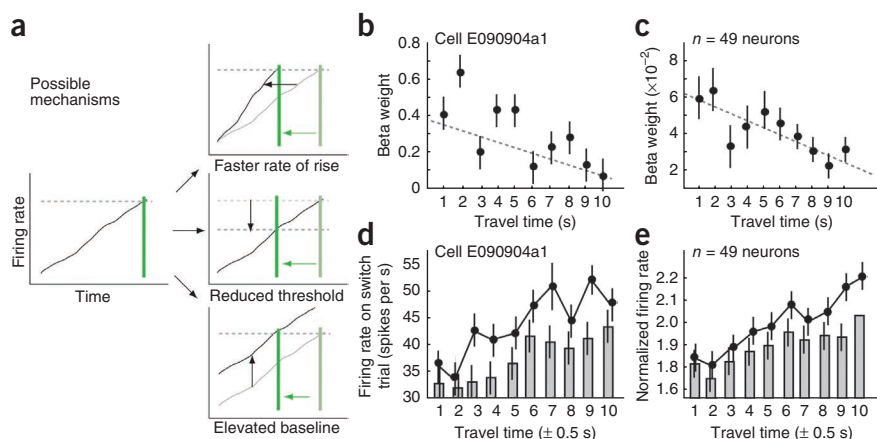
We next calculated the average neural responses in a series of 5-s time bins encompassing multiple choices separated by earliness level. For both the example neuron and the focal population of neurons ( $n = 49$  of 102, see above), we found that the rate of rise of firing rates was positively correlated with earliness (**Fig. 4c,d**). We quantified these effects by calculating the regression weight for firing rate as a function of patch residence time separately for each of the four earliness bins. The slope for the earliest patches ( $\beta = 0.71$ ) was greater than the slope for the early patches ( $\beta = 0.52$ ,  $P < 0.01$ , bootstrap test), which was greater than the slope for the late patches ( $\beta = 0.44$ ,  $P < 0.01$ ), which was in turn greater than the slope for the latest patches ( $\beta = 0.39$ ,  $P < 0.01$ ). The same effects were observed in the population average responses ( $P < 0.005$  for each comparison). Slopes decreased monotonically from earliest to latest quartiles for 32 of 49 neurons (65%). The remainder ( $n = 17$ ) showed no such effect.

We then examined whether firing rates rose to the same threshold in each of the four earliness bins. For this analysis, we only examined firing rates for patch-leaving choices, which we had ignored in previous analyses. We took these firing rates to be a proxy for the patch-leaving threshold. Firing rates did not depend on earliness for the example neuron (regression of firing rate against earliness level,  $P = 0.45$ ; **Fig. 4c**) or the population ( $P = 0.88$ ; **Fig. 4d**), which is consistent with the threshold hypothesis. We observed no relationship between earliness and threshold for the population average response (regression,  $P = 0.79$ ). We observed a significant effect of earliness on threshold for only a small number of cells, 6 out of 49 significantly modulated neurons (12.2%,  $P < 0.05$ ). Finally, we found that neuronal responses aligned to the patch-leaving trial for different earliness levels overlapped (**Fig. 4e**). For the population of neurons, firing rates rose to approximately the same level on the last decision to stay in the patch before choosing to leave the patch (regression,  $\beta = -0.003$ ,  $P = 0.51$ ).

### Travel time influences gain and threshold of neuronal responses

Assuming that patch-leaving decisions are governed by a threshold process, variation in travel times should influence the threshold. There are three basic mechanisms by which travel time could influence the accumulation-to-threshold process (**Fig. 5a**). First, travel time could increase the rate of rise of the decision variable. Second, travel time could adjust the threshold level. Third, travel time could influence the baseline. Our next analysis was designed to determine which of these processes are implemented in dACC.

We first examined the relationship between travel time and response gain (**Fig. 5b,c**). For each neuron, we divided patches into ten travel time deciles (equal-sized sequentially classified bins, 1 to  $10 \pm 0.5$  s). In each decile, we calculated the slope of firing rate versus time in patch. We found a significant negative correlation between travel time and regression slope for our example neuron (regression,  $\beta = -0.04$ ,  $P < 0.01$ ; **Fig. 5b**) and for the 49 cells in the analysis population ( $\beta = -0.033$ ,  $P < 0.01$ ; **Fig. 5c**). We observed significant negative effects in 29 of 49 of the focal population of neurons, positive correlations in 6, and no effect in 14 ( $P < 0.05$ ). Across the entire population, we found a significant negative correlation in 41 of 102 neurons ( $P > 0.05$ ), a positive correlation in 7 neurons and no effect in 54 neurons.



**Figure 5** Travel time governs both neuronal response gain and threshold. **(a)** Schematic of three possible mechanisms by which exogenous factors may govern a rise-to-threshold process. Shorter travel times can hasten patch-leaving (leftward movement on x axis) by increasing the rate of rise, reducing the threshold or elevating the baseline. **(b,c)** Evidence that travel times change rate of rise. Example neuron **(b)** and population **(c)** average regression slopes (beta weights) for firing rate as a function of time in patch. Beta weights fell as travel times rose, indicating that shorter travel time increases neuronal response gain. Error bars represent s.e.m. **(d,e)** Evidence that travel time influences firing threshold for patch abandonment. Firing rate on patch-leaving trial was taken as a proxy for threshold level. Example neuron **(d)** and population **(e)** show increasing firing rates on patch-leaving trial as travel time increases (black dots). Firing rate on penultimate trial also rose with travel time, consistent with a multi-trial integration process. Error bars represent s.e.m.

We next examined the relationship between travel time and threshold (Fig. 5d,e). As above, we assumed that the firing rate in the epoch immediately preceding patch-leaving provides a proxy for threshold. For our example neuron, the threshold rose with travel time ( $\beta = 1.53$ ,  $P < 0.01$ , regression of firing rates (spikes per s) against travel time (s); Fig. 5d). There was also a significant, but weaker, correlation between travel time and firing rate on the last choice to stay in the patch before patch-leaving, consistent with the idea of a gradual rise-to-threshold process ( $\beta = 1.32$ ,  $P < 0.025$ ); the whole population showed similar effects (Fig. 5e). For the focal population ( $n = 49$ ), threshold rose with travel time ( $\beta = 0.92$ ,  $P < 0.01$ ), as did previous-trial responses ( $\beta = 0.69$ ,  $P < 0.001$ ). We observed a rise in threshold in 37 of 102 neurons (regression,  $P < 0.05$ , a negative effect in one neuron, no effect in the rest).

If travel time influences threshold levels, do thresholds actually remain constant for different leaving times, after accounting for travel time? To answer this question, we measured neural responses on patch-leaving trials in five groups of travel times (0.5 to 2.5, 2.5 to 4.5, 4.5 to 6.5, 6.5 to 8.5 and 8.5 to 10.5 s). Two facts suggest that they were constant across earliness level (see above) after accounting for travel time. First, in each of these travel time groups, earliness did not influence firing rate in the pre-saccadic epoch of the patch-leaving trial ( $P > 0.2$  in all cases). Second, there was no significant interaction between earliness and travel time ( $P = 0.67$ , two-way ANOVA, four levels of earliness and five travel times, main effect of earliness = 0.06 spikes per s per bin).

Finally, we considered whether travel times influence baseline firing. We reasoned that such an effect would be most apparent on the first few choices in a new patch. We found no effects of travel time on firing rate responses occurring in the first 5 s of a patch (ANOVA of normalized firing rate against travel time,  $P = 0.41$ ) and a significant effect in nine neurons individually ( $P < 0.05$ ). The frequency of these effects is not much more than would be expected by chance ( $n = 5.1$  neurons,  $P = 0.067$ , binomial test), and the size of the effect was weak (average

difference between 1- and 10-s travel times was 0.09 spikes per s). Moreover, of these nine neurons, six showed increasing firing rates with longer travel times. Our data therefore do not endorse the idea that travel times influence baseline neuronal firing rates in ACC.

## DISCUSSION

Choosing when to leave a depleting resource patch is a ubiquitous natural decision problem that is central to foraging theory and behavioral ecology<sup>2,3</sup>. Although the brains of animals have undoubtedly been shaped by evolutionary pressures for foraging efficiency, the neural processes that mediate the simple decision to give up on one patch and move to another remain obscure. Our findings suggest that, during foraging, the primate brain computes a decision variable whose magnitude corresponds to the relative value of leaving a patch, that this value rises to a threshold associated with patch leaving, and that travel time between patches governs both the threshold and the rate at which this decision variable rises. This decision variable is represented in the firing rates of neurons in

the dACC, a frontal lobe structure associated with reward monitoring and behavioral adjustment<sup>14–16</sup>.

If the thresholding process that we propose were to occur at the level of dACC or its inputs, we would expect to see the outcome of the threshold process in the responses of dACC neurons. Instead, we observed a signal that varied continuously with time in patch, suggesting that the thresholding process occurs downstream of dACC, perhaps in FEFs or some other premotor structure. Notably, our results imply that a downstream neuron cannot judge whether it is time to leave the patch solely by querying the output of dACC; it needs to have information about travel time, likely in the form of the change in its threshold. We speculate that such control may be implemented through neuromodulatory inputs to the dACC, perhaps via dopamine or norepinephrine<sup>36,37</sup>.

Responses of dACC neurons do not uniquely represent any single variable in the MVT equations. The MVT is a description of foraging behavior at the computational level, whereas our data support a particular mechanism by which the decision process could be implemented<sup>38</sup>. Thus, although we claim that our hypothesized decision variable encodes the relative value of leaving a patch, we could just as easily argue that it encodes the negative value of staying (indeed, this would be consistent with error-related theories of ACC function<sup>16</sup>). Because these functional mechanisms resemble those known to support basic perceptual and mnemonic decisions<sup>10–13,34</sup>, our findings endorse the idea that the brain uses a small suite of common mechanisms to solve diverse problems in multiple domains.

The broad applicability of the MVT to such a wide array of organisms underscores the fact that dACC is unlikely to be the sole neural locus of the decision process, even in organisms with brains similar to ours. Indeed, these mechanisms may not even be limited to brains. Many organisms that lack brains, including amoebas, slime molds and plants, exhibit behavior that is consistent with the MVT<sup>6,39,40</sup>. We conjecture that such organisms solve the patchy foraging problem in much the same way that monkeys do, namely by maintaining and

controlling a representation of the relative value of leaving a patch. Thus, diverse organisms may solve common problems using similar algorithms that are implemented very differently<sup>38</sup>. Even in organisms with brains similar to ours, given the high redundancy of decision signals across brain areas in primates<sup>41</sup>, we predict that similar signals might be observed in other regions, including the dorsolateral prefrontal cortex, lateral intraparietal area and posterior cingulate cortex, although such signals might be convolved with other information such as target location or movement metrics.

### Relation to previous studies

Our findings are broadly consistent with prior studies showing that dACC monitors reward information from many sources and signals the need to adjust behavior in some manner<sup>17,18,22,23,42–46</sup>. Our results corroborate these earlier results and extend them in four important ways. First, we found that dACC responses vary continuously with the extent to which circumstances favor the decision to move on, even if leaving does not occur. This observation supports the idea that dACC neurons represent a scalar decision variable reflecting the relative value of leaving. The relative value of switching behavior was not manipulated in a previous study in which all non-switch trials were, in all important respects, the same<sup>23</sup>. Second, we found and measured a specific threshold at which leaving, and by extension switching, occurs. Third, we identified two mechanisms by which exogenous factors govern patch-leaving behavior. Finally, we found that neuronal activity in dACC promotes disengagement in a relatively natural task that is directly modeled on real-world foraging situations. These results directly link dACC neuronal activity to behavior in a naturalistic context and may extend to situations outside the laboratory in which dACC dysfunction has been implicated, including addiction, depression, obsessive-compulsive disorder and Tourette syndrome<sup>25–27</sup>.

At first glance, our results appear to contradict those obtained by an earlier study that reported increasing firing rates of dACC neurons with increasing proximity to reward<sup>24</sup>, whereas we found increasing firing rates in anticipation of sequentially smaller rewards. We believe that the two sets of findings are fully concordant. In our study, firing rates rose as the monkey approached the decision to abandon the current patch for a new one. In the earlier study, firing rates increased as the monkey neared the rewarded action. In both cases, firing rates of dACC neurons marked progression through a sequence of actions toward a salient behavioral event: the reward in their task, patch-leaving in ours. Together, our results suggest a broader view, namely that dACC neurons do not signal reward value *per se*, but rather that their responses encode an abstract decision variable that is suitable for guiding a variety of different modifications in behavior, whether generated endogenously or exogenously.

Our findings may also initially appear to contradict results from our earlier studies of dACC neurons<sup>21,22</sup>. Previously, we found that the firing rates of dACC neurons reflected both real and fictive rewards, and generally did so with higher firing rates for larger rewards<sup>22</sup>. In that study, however, large rewards, both real and fictive, promoted a behavioral strategy that led to potentially larger rewards. Indeed, trial-to-trial variations in firing rate in that study positively covaried with likelihood of adjusting behavior on the next trial. In other words, higher firing rates in both studies predicted the likelihood that the monkey would successfully incorporate new information about the world into an ongoing decision to change behavior.

In another study, we found that neuronal activity in dACC showed weak, but significant, selectivity for saccade direction, in addition to anticipated real and fictive reward size<sup>21</sup>. In contrast, here we found no evidence for spatial selectivity in neuronal responses. These differ-

ences likely reflect task design. We used eight targets in the previous study, but only two saccade targets here, thus weakening our sensitivity to spatial selectivity, especially bimodal tuning common in dACC<sup>21</sup>. Also, the task used in the prior study demanded that monkeys carefully distinguish adjacent, physically similar targets to evaluate the associated reward outcomes, whereas the two targets were widely separated and physically distinct in the current study. We hypothesize that the greater demand for attentional resources associated with spatial locations in that task accentuated spatial tuning in the earlier study.

### Conclusion

Our virtual foraging task is an ersatz idealization of a real patchy foraging environment. Given that foraging often involves physical effort, these results are only a first step on the path to understanding real foraging decisions. Because of the clear links between ACC function and effortful choices, dACC seems particularly well positioned to guide real-world foraging choices and is likely involved in these choices. Thus, we believe that our results provide a useful advance toward understanding natural value-based decisions and forge a critical link between systems neurobiology and behavioral ecology.

Animals' bodies have evolved to efficiently exploit the resources in their environments. Natural selection has also acted on the nervous systems of these animals to enable the adaptive action of their bodies. Few studies have linked neural computations to specific types of naturally occurring foraging decisions. Our study portends a more general understanding of prey selection, diet selection and more complex foraging problems<sup>3,47</sup>. Ultimately, these results endorse the unification of theoretical and experimental work in the ecological and neural sciences<sup>48</sup>.

### METHODS

Methods and any associated references are available in the online version of the paper at <http://www.nature.com/natureneuroscience/>.

*Note: Supplementary information is available on the Nature Neuroscience website.*

### ACKNOWLEDGMENTS

We thank S. Heilbronner for comments on design, analysis and writing. This research was supported by US National Institutes of Health grant R01EY013496 (M.L.P.), a Fellowship from the Tourette Syndrome Association (B.Y.H.) and US National Institutes of Health grant K99 DA027718-01 (B.Y.H.).

### AUTHOR CONTRIBUTIONS

B.Y.H. designed the experiment and collected the data. B.Y.H. and J.M.P. contributed to data analysis. B.Y.H., J.M.P. and M.L.P. wrote the manuscript.

### COMPETING FINANCIAL INTERESTS

The authors declare no competing financial interests.

Published online at <http://www.nature.com/natureneuroscience/>.

Reprints and permissions information is available online at <http://www.nature.com/reprints/index.html>.

1. Prins, H.H.T. *Ecology and Behavior of the African Buffalo: Social Inequality and Decision Making* (Chapman and Hall, London, 1996).
2. Charnov, E.L. Optimal foraging, the marginal value theorem. *Theor. Popul. Biol.* **9**, 129–136 (1976).
3. Stephens, D.W. & Krebs, J.R. *Foraging Theory* (Princeton University Press, Princeton, NJ, 1986).
4. Bendesky, A., Tsunozaki, M., Rockman, M.V., Kruglyak, L. & Bargmann, C.I. Catecholamine receptor polymorphisms affect decision-making in *C. elegans*. *Nature* **472**, 313–318 (2011).
5. Thompson, D. & Fedak, M.A. How long should a dive last? A simple model of foraging decisions by breath-hold divers in a patchy environment. *Anim. Behav.* **61**, 287–296 (2001).
6. McNickle, G.G. & Cahill, J.F. Jr. Plant root growth and the marginal value theorem. *Proc. Natl. Acad. Sci. USA* **106**, 4747–4751 (2009).
7. Smith, E.A. & Winterhalder, B. *Evolutionary Ecology and Human Behavior* (de Gruyter, New York, 1992).

8. Agetsuma, N. Simulation of patch use by monkeys using operant conditioning. *J. Ethology* **16**, 49–55 (1999).
9. Gold, J.I. & Shadlen, M.N. The neural basis of decision making. *Annu. Rev. Neurosci.* **30**, 535–574 (2007).
10. Gold, J.I. & Shadlen, M.N. Banburismus and the brain: decoding the relationship between sensory stimuli, decisions, and reward. *Neuron* **36**, 299–308 (2002).
11. Hanes, D.P. & Schall, J.D. Neural control of voluntary movement initiation. *Science* **274**, 427–430 (1996).
12. Schall, J.D. On building a bridge between brain and behavior. *Annu. Rev. Psychol.* **55**, 23–50 (2004).
13. Carpenter, R.H.S. *Movements of the Eyes* (Pion, London, 1988).
14. Kennerley, S.W., Walton, M.E., Behrens, T.E., Buckley, M.J. & Rushworth, M.F. Optimal decision making and the anterior cingulate cortex. *Nat. Neurosci.* **9**, 940–947 (2006).
15. Rushworth, M.F. & Behrens, T.E. Choice, uncertainty and value in prefrontal and cingulate cortex. *Nat. Neurosci.* **11**, 389–397 (2008).
16. Holroyd, C.B. & Coles, M.G. The neural basis of human error processing: reinforcement learning, dopamine, and the error-related negativity. *Psychol. Rev.* **109**, 679–709 (2002).
17. Seo, H. & Lee, D. Temporal filtering of reward signals in the dorsal anterior cingulate cortex during a mixed-strategy game. *J. Neurosci.* **27**, 8366–8377 (2007).
18. Quilodran, R., Rothe, M. & Procyk, E. Behavioral shifts and action valuation in the anterior cingulate cortex. *Neuron* **57**, 314–325 (2008).
19. Ito, S., Stuphorn, V., Brown, J.W. & Schall, J.D. Performance monitoring by the anterior cingulate cortex during saccade countermanding. *Science* **302**, 120–122 (2003).
20. Williams, Z.M., Bush, G., Rauch, S.L., Cosgrove, G.R. & Eskandar, E.N. Human anterior cingulate neurons and the integration of monetary reward with motor responses. *Nat. Neurosci.* **7**, 1370–1375 (2004).
21. Hayden, B.Y. & Platt, M.L. Neurons in anterior cingulate cortex multiplex information about reward and action. *J. Neurosci.* **30**, 3339–3346 (2010).
22. Hayden, B.Y., Pearson, J.M. & Platt, M.L. Fictive reward signals in the anterior cingulate cortex. *Science* **324**, 948–950 (2009).
23. Shima, K. & Tanji, J. Role for cingulate motor area cells in voluntary movement selection based on reward. *Science* **282**, 1335–1338 (1998).
24. Shidara, M. & Richmond, B.J. Anterior cingulate: single neuronal signals related to degree of reward expectancy. *Science* **296**, 1709–1711 (2002).
25. Baler, R.D. & Volkow, N.D. Drug addiction: the neurobiology of disrupted self-control. *Trends Mol. Med.* **12**, 559–566 (2006).
26. Stern, E. *et al.* A functional neuroanatomy of tics in Tourette syndrome. *Arch. Gen. Psychiatry* **57**, 741–748 (2000).
27. Devinsky, O., Morrell, M.J. & Vogt, B.A. Contributions of anterior cingulate cortex to behavior. *Brain* **118**, 279–306 (1995).
28. Stephens, D.W. & Anderson, D. The adaptive value of preference for immediacy: when shortsighted rules have farsighted consequences. *Behav. Ecol.* **12**, 330–339 (2001).
29. Evans, T.A. & Beran, M.J. Delay of gratification and delay maintenance by rhesus macaques (*Macaca mulatta*). *J. Gen. Psychol.* **134**, 199–216 (2007).
30. Kim, S., Hwang, J. & Lee, D. Prefrontal coding of temporally discounted values during intertemporal choice. *Neuron* **59**, 161–172 (2008).
31. Louie, K. & Glimcher, P.W. Separating value from choice: delay discounting activity in the lateral intraparietal area. *J. Neurosci.* **30**, 5498–5507 (2010).
32. Kahneman, D., Knetsch, J.L. & Thaler, R.H. Anomalies: the endowment effect, loss aversion, and the status quo bias. *J. Econ. Perspect.* **5**, 193–206 (1991).
33. Mazur, J.E. An adjusting procedure for studying delayed reinforcement. in *Quantitative Analyses of Behavior, vol 5. The Effect of Delay and Intervening Events on Reinforcement Value* (eds. Commons, M.L., Mazur, J.E., Nevin, J.A. & Rachlin, H.) (Erlbaum, Mahway, New Jersey, 1987).
34. Roitman, J.D. & Shadlen, M.N. Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *J. Neurosci.* **22**, 9475–9489 (2002).
35. Horwitz, G.D., Batista, A.P. & Newsome, W.T. Representation of an abstract perceptual decision in macaque superior colliculus. *J. Neurophysiol.* **91**, 2281–2296 (2004).
36. Tobler, P.N., Fiorillo, C.D. & Schultz, W. Adaptive coding of reward value by dopamine neurons. *Science* **307**, 1642–1645 (2005).
37. Yu, A.J. & Dayan, P. Uncertainty, neuromodulation, and attention. *Neuron* **46**, 681–692 (2005).
38. Marr, D.C. *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information* (Freeman, New York, 1982).
39. Latty, T. & Beekman, M. Food quality affects search strategy in the acellular slime mold, *Physarum polycephalum*. *Behav. Ecol.* **20**, 1160–1167 (2009).
40. Bonser, R., Wright, P.J., Bament, S. & Chukwu, U.O. Optimal patch use by foraging workers of *Lasius fuliginosus*, *L. niger* and *Myrmica ruginodis*. *Ecol. Entomol.* **23**, 15–21 (1998).
41. Wallis, J.D. & Kennerley, S.W. Heterogeneous reward signals in prefrontal cortex. *Curr. Opin. Neurobiol.* **20**, 191–198 (2010).
42. Sallet, J. *et al.* Expectations, gains, and losses in the anterior cingulate cortex. *Cogn. Affect. Behav. Neurosci.* **7**, 327–336 (2007).
43. Kennerley, S.W., Dahmubed, A.F., Lara, A.H. & Wallis, J.D. Neurons in the frontal lobe encode the value of multiple decision variables. *J. Cogn. Neurosci.* **21**, 1162–1178 (2008).
44. Matsumoto, M., Matsumoto, K., Abe, H. & Tanaka, K. Medial prefrontal cell activity signaling prediction errors of action values. *Nat. Neurosci.* **10**, 647–656 (2007).
45. Amiez, C., Joseph, J.P. & Procyk, E. Reward encoding in the monkey anterior cingulate cortex. *Cereb. Cortex* **16**, 1040–1055 (2006).
46. Hayden, B.Y., Heilbronner, S.R., Pearson, J.M. & Platt, M.L. Surprise signals in anterior cingulate cortex: neuronal encoding of unsigned reward prediction errors driving adjustment in behavior. *J. Neurosci.* **31**, 4178–4187 (2011).
47. Stephens, D.W., Brown, J.S. & Ydenberg, R.C. *Foraging: Behavior and Ecology* (University of Chicago Press, Chicago, 2007).
48. Wilson, E.O. *Consilience: The Unity of Knowledge* (Knopf, New York, 1998).

## ONLINE METHODS

**Surgical procedures.** All procedures were approved by the Duke University Institutional Animal Care and Use Committee and were designed and conducted in compliance with the Public Health Service's Guide for the Care and Use of Animals. Two male rhesus monkeys (*Macaca mulatta*) served as subjects. Initially, a head-holding prosthesis was implanted in both animals using standard surgical techniques. Six weeks later, animals were habituated to laboratory conditions and trained to perform oculomotor tasks for liquid reward. A second surgical procedure was then performed to place a plastic recording chamber (Crist Instruments) over dorsal anterior cingulate cortex. Animals received analgesics and antibiotics after all surgeries. Throughout both behavioral and physiological recording sessions, the chamber was kept sterile with regular antibiotic washes and sealed with sterile Teflon caps.

**Behavioral task.** Monkeys were placed on controlled access to fluid outside of experimental sessions to motivate behavior. Horizontal and vertical eye positions were sampled at 1,000 Hz by an infrared eye-monitoring camera system (SR Research). Stimuli were controlled by a computer running MATLAB (MathWorks) with Psychtoolbox and Eyelink Toolbox<sup>49,50</sup>. Visual stimuli were small colored rectangles on a computer monitor placed directly in front of the animal and centered on his eyes. A standard solenoid valve controlled the duration of juice delivery. We estimated the precision of fluid volume delivered by the solenoid across the range of open time commands used in this study. Given that the s.e.m. in volume increased with the mean, we calculated the coefficient of variation as a measure of precision, and this value was 0.29. Because of this small uncertainty in liquid volume delivered, the reported values should be taken as approximate.

Our task was designed to mimic the critical elements of the patch-leaving task studied in foraging theory<sup>2,3</sup>. On each trial, a small yellow square appeared in the center of the monitor (<0.5 degrees of visual angle). Following its appearance, the monkey aligned gaze with the square to indicate his readiness to begin the trial. Once the monkey acquired fixation ( $\pm 0.5$  degrees), two eccentric targets appeared, one on the left and one on the right (300 pixels from fixation). One target was a small blue rectangle (remain in patch option) and the other was a gray rectangle (leave patch option) that was the same width (80 pixels) and taller (precise height depending on condition). One pixel is approximately equal to 0.024 degrees of visual angle.

Following a 500-ms delay, the central fixation square extinguished and the monkey was free to select either of the two targets by shifting gaze to it ( $\pm 2$  degrees from the center of the rectangle). Following choice of either target, the rectangle began to shrink at a constant rate (65 pixels per s) until it disappeared, a reward was given (if the blue 'stay' target was chosen), and the ITI began (1 s). Because the rate of bar shrinking was constant, the height of the bar provided an unambiguous cue to the delays associated with the two options on every trial.

The delay associated with the blue stay (that is, remain in patch) rectangle occurred before the reward and was isomorphic to the handling time in foraging decisions (set at 400 ms in the task used for recording, that is, the variable travel time task). The delay associated with the gray 'switch' (that is, leave the patch) rectangle was analogous to the travel time in foraging decisions (ranging from 0.5 to 10.5 s in this experiment). It was set at a random value on each patch, but did not vary in a patch. The fixed delay (ITI) between trials was uncued, but was always the same (1 s). This fixed delay is isomorphic to a fixed search time in foraging theory. Errors (14% of trials) consisted of either early fixation breaks during the brief hold period (89% of the error trials) or failures to initiate the trial within 30 s (the remaining 11% of error trials). In either case, an error was followed by the presentation of a dark green square on the center of the screen (an error cue) and a 3-s timeout period. We defined search time as the average total time per trial in each patch, excluding handling time. Search time included the ITI, the saccade time, and any other sources of delay or variability.

Following the first choice of the blue stay rectangle in each patch, the monkey received 306  $\mu$ l of water. On subsequent choices of the 'stay' target, the reward decreased by 19  $\mu$ l (although we introduced a small variance in this amount,  $\epsilon$  = s.e.m. of 1.9  $\mu$ l). If the monkey continued to choose the blue stay option, its value would eventually reach 0 and remain 0 thereafter. On choosing the gray switch rectangle, the location of the blue and gray rectangles would alternate and the value of the blue rectangle would reset to 306  $\mu$ l. On choosing the gray rectangle, the size of the gray rectangle and the associated travel time would reset to a new value, chosen from a uniform distribution between 0.5 and 10.5 s.

Rewards were not completely deterministic in this task, although the variability was quite small. Stochasticity came from two sources. First, there was an inherent variability in the reward amounts, added to keep the task engaging, and to encourage the monkeys to actively monitor rewards, rather than simply learn to count a certain number of trials. Second, there is necessarily some uncontrolled, but small, variability associated with the juice dispenser. We estimate that this variability is less than that generated by the variability that we deliberately added.

To plot residence time as a function of travel time (Fig. 2), we computed residence time for each patch. Residence time includes all time in patch from arrival to the decision to leave, and thus included handling time and search time and its constituents, saccade latencies, reward duration, etc. On the occasional trial in which the monkey chose to leave the patch immediately on the first trial, we assigned the residence time a value of 0 s.

The critical equation in the MVT is equation (2) from ref. 3.

$$E_n = \frac{\sum P_i \times g_i(T_i) - t \times E_T}{t + \sum P_i \times T_i}$$

In this equation,  $E_n$  is net energy intake (the quantity to be maximized),  $P_i$  is the proportion of patches to be visited of a given type,  $g_i$  is the assimilated energy corrected for the cost of searching (assumed to be zero in this task),  $T$  is the hunting time (in this case, patch residence time),  $t$  is the travel time, and  $E_T$  is the energy cost per unit time for traveling.

We adopted a simplified version of the more general situation encompassed by the MVT. Two simplifying assumptions are especially important. First, all patches are identical. Second, there is no energy cost associated with travel, only a time cost.

Thus, intake rate for a given patch is given by  $E_n = \frac{g(T)}{t + T}$ . According to the

MVT, intake rate is maximized when  $E_n^* = \frac{\delta g(T)}{\delta T}$ , where  $E_n^*$  refers to reward intake rate when patch residence time is maximized. At this point, the marginal intake rate matches the average intake rate for the habitat.

Total patch residence time depends strongly on several factors that are beyond the control of the monkeys. These include handling time and ITI, but also the monkey's own deliberation time, reaction times, saccade latency, etc. As these variables potentially lie outside the control of the monkeys, we assumed that they treat them as fixed quantities. Reward as a function of average patch residence time is shown in Figure 1b. Because the reward function was somewhat stochastic, we used a repeated randomized algorithm to calculate rate-maximizing behavior. Specifically, we simulated the total reward harvested over 5,000 trials for each of 100 patch residence times, ranging from 0 to 75 s. To reduce any possible effects of noise, we repeated each simulation 1,000 times. We determined the rate-maximizing value by locating the peak of the resulting intake curve. This randomization process eliminates biases emanating from the variability in reward values.

**Microelectrode recording techniques.** Single electrodes (Frederick Haer) were lowered using a microdrive (Kopf) until the waveform of a single (1–3) neuron(s) was isolated. Individual action potentials were identified by standard criteria and isolated on a Plexon system (Plexon). Neurons were selected for study on the basis of the quality of isolation, and not on task-related response properties.

We approached dACC through a standard Teflon recording grid (Crist Instruments). dACC was identified by structural magnetic resonance images taken before the experiment. Neuroimaging was performed at the Center for Advanced Magnetic Development at DUMC, on a 3T Siemens Medical Systems Trio MR Imaging Instrument using 1-mm slices. We confirmed that electrodes were in dACC using stereotactic measurements, as well as by listening for characteristic sounds of white and gray matter during recording. Our recordings were likely to have come from area 24, and especially the dorsal and ventral banks of the anterior cingulate sulcus (see Supplementary Fig. 8 for reconstructions of recording sites).

Prior to beginning formal experiments, we performed several exploratory recording sessions to map the physiological response properties of neurons accessible through our recording chamber. We were able to distinguish white from gray matter by the presence and absence of neural activity. During these mapping sessions, we were able to identify both the dorsoventral and mediolateral

extent of the cingulate sulcus. Neurons were recorded in the same animals and at the same grid positions as used in two previous studies examining the role of dACC in representing information about rewards<sup>22</sup> and action<sup>21</sup>. It is therefore likely that neurons in both studies come from the same brain region. Moreover, as in these previous studies, we performed no pre-selection on neurons, aside from the natural biases in single-unit recordings toward larger and higher firing rate neurons<sup>21</sup>.

**Analysis.** PSTHs were constructed by aligning spike rasters to trial events and averaging firing rates across multiple trials. Firing rates were calculated in 10-ms bins,

but were generally analyzed in longer epochs. For display, PSTHs were smoothed using a 100-ms running boxcar. Neuronal activity was temporally aligned to the end of the saccade indicating the animal's choice (time zero). To normalize neuronal activity, we calculated the average firing rate for each neuron during a pre-trial epoch, defined as a 0.5-s period beginning 0.75 s before the beginning of each trial. We then divided neural activity by this value for each neuron.

49. Brainard, D.H. The Psychophysics Toolbox. *Spat. Vis.* **10**, 433–436 (1997).

50. Cornelissen, F.W., Peters, E. & Palmer, J. The EyeLink Toolbox: eye tracking with MATLAB and the Psychophysics Toolbox. *Behav. Res. Methods Instrum. Comput.* **34**, 613–617 (2002).