

Neurons in the Rat Anterior Cingulate Cortex Dynamically Encode Cost–Benefit in a Spatial Decision-Making Task

Kristin L. Hillman and David K. Bilkey

Department of Psychology, University of Otago, Dunedin 9054, New Zealand

Optimal decision-making often requires an assessment of the costs and benefits associated with each available course of action. Previous studies have shown that lesions to the anterior cingulate cortex (ACC) impair cost–benefit decision-making in laboratory animals, but the neural mechanisms underlying the deficit are not well understood. We recorded from ACC neurons in freely moving rats as they performed a spatial decision-making task whereby, in the baseline configuration “2:6B,” rats could pursue two or six food pellets, the latter obtained by climbing a barrier [high cost, high reward (HCHR)]. In this configuration, the mean percentage of HCHR choices was $69 \pm 4\%$, and a substantial portion of ACC neurons (63%) exhibited significantly higher firing for one goal trajectory versus the other; for 94% of these cells, higher firing was associated with the HCHR option. This HCHR bias was not simply attributable to the larger reward, the barrier, or behavioral preference. In intersession and intrasession manipulations involving at least one barrier (2:6B, 2B:6B, and 2:2B), ACC activity rapidly adapted and was consistently biased toward the economically advantageous option relative to the configuration. Interestingly, when only a difference in reward magnitude was presented (2:6, no barrier, HCHR choices of $84 \pm 4\%$), ACC activity was minimal and nonbiased. One interpretation of our data is that the ACC encodes a relative, integrated cost–benefit representation of available choice options that is biased toward the “better” option in terms of effort/outcome ratio. This representation may be specifically recruited when an assessment of reward and effort is required to optimally perform a task.

Introduction

In the past decade, there has been a growing interest in the role of the anterior cingulate cortex (ACC) in choice behavior. The ACC is anatomically well positioned to integrate reward information with action selection given its cortico-cortico, sensorimotor, and subcortical connections (Carmichael and Price, 1995a,b; Hoover and Vertes, 2007), and imaging shows that the region is activated during voluntary decision-making (Walton et al., 2004; Mars et al., 2005; Forstmann et al., 2006; Yoshida and Ishii, 2006; Croxson et al., 2009). Lesions to the ACC are associated with suboptimal choice behavior. ACC lesioned primates fail to use feedback, either positive or negative, to optimize their behavior during choice tasks (Amiez et al., 2006; Kennerley et al., 2006), and transient inactivation of the region produces perseverance in behaviors that are no longer beneficial (Shima and Tanji, 1998). ACC-lesioned rats show a profound deficit in investing physical effort to reach larger food rewards, a course of action they consistently elect preoperatively to maximize reward gain (Walton et al., 2003; Schweimer et al., 2005; Floresco and Ghods-Sharifi, 2007).

One interpretation of available data is that the ACC biases behavior toward actions that will maximize overall reward gain, possibly by encoding a cumulative history of recent reward (Seo and Lee, 2007; Walton and Mars, 2007). As a result,

ACC-lesioned animals might act at chance levels in choice tasks because they lack a representation of which choice has the highest rate of return. Although this view is supported by data obtained from monkeys, it does not explain why ACC-lesioned rats shift their behavior away from high-cost, high-reward (HCHR) options despite being capable of discriminating reward size (Walton et al., 2002; Rudebeck et al., 2006) and reward location (Kesner and Ragozzino, 2003). This suggests that the ACC must also be sensitive to the effort requirements of a task. Studies of single-unit firing in the ACC have confirmed that neurons in this region respond to independent variations in task effort (Kennerley et al., 2009) or reward size (Amiez et al., 2006; Sallet et al., 2007; Kennerley et al., 2009), but it is not clear whether this information is represented in any integrated manner in this region. The purpose of the present study is to address this question.

Here we recorded from neurons in the ACC of freely moving rats as the animals explored two options that differed in cost/benefit ratio. Cost related to the degree of physical effort required to obtain food reward. We hypothesized that the function of the ACC is to encode a cost–benefit representation of each course of action currently available to the animal. We predicted that such a representation would be (1) reflective of reward magnitudes, discounted by the effort required to obtain them, (2) indicative of the relative cost–benefit value based on the available alternatives, and (3) adaptive to external changes in the cost–benefit environment. Such a representation could be activated during decision-making to enable optimal choice behavior, particularly in situations in which investing cognitive or physical effort will merit larger rewards.

Received March 11, 2010; revised April 13, 2010; accepted April 27, 2010.

This work was supported by The Royal Society of New Zealand Marsden Fund.

Correspondence should be addressed to David K. Bilkey, Department of Psychology, University of Otago, P.O. Box 56, Dunedin 9054, New Zealand. E-mail: dbilkey@psy.otago.ac.nz.

DOI:10.1523/JNEUROSCI.1273-10.2010

Copyright © 2010 the authors 0270-6474/10/307705-09\$15.00/0

Materials and Methods

Subjects. Ten male Sprague Dawley rats (Hercus-Taieri Resource Unit) weighing 400–680 g were used in the study. Rats were single housed in translucent plastic cages containing pine chips and maintained on a 12 h light/dark cycle. All training and experimentation occurred during the light phase. After 2 weeks of daily handling and weighing, animals were food deprived of standard rat chow (Specialty Feeds) to no less than 85% of their free-feeding weight to promote interest in food reward (Coco Pop cereal; Kellogg's) during test phases. Water was available *ad libitum* in the home cage.

Preoperative training. In the initial week of training, rats were individually habituated for 15 min/d to the experimental setup, a continuous T-maze (Fig. 1*a*), 90 × 80 × 22 cm constructed of particle board and painted black. Cereal pellets were scattered throughout the maze to promote exploration. In the second week of training, four cereal pellets were placed at each reward site, and rats were trained to run the maze in a unidirectional manner. Starting at the bottom of the midstem, each rat learned to run up the midstem, turn right or left, and proceed around to a baited reward area. After consuming the cereal reward, the rat continued its unidirectional path around the base of the maze and then turned into the midstem to initiate another trial. Reversals in unidirectional travel and attempts to circumnavigate the midstem (i.e., travel straight across the base toward the other food reward location) were blocked by the experimenter with a particle board insert. The rat was not paused between trials but completed trials in a continuous, uninterrupted manner at a pace which the rat itself set. In the third and fourth weeks of training, four cereal pellets were placed at each reward site, and wedge-shaped barriers with a right-angle profile were gradually introduced to both arms of the maze. The barriers were constructed of cardboard with a wire grid overlay for added traction. Rats had to surmount the vertical face of the wedge and travel down the sloping face to reach the reward site. Barrier height was initially 15 cm and increased 5 cm every third day until the full height of 30 cm was achieved.

Surgery. Once animals were proficient at running and barrier climbing, they were surgically implanted with an adjustable microdrive assembly mounted on a McIntyre miniconnector head plug as described previously (Bilkey et al., 2003). Briefly, animals were anesthetized by intraperitoneal injection of ketamine (75 mg/kg) and domitor (0.5 mg/kg). Seven 25 μ m Formvar-coated nichrome wires (California Fine Wire) were inserted into the ACC via stereotaxic guided craniotomy, with coordinates +1.7 mm anteroposterior and \pm 0.4 mm mediolateral from bregma, –1.8 mm dorsoventral from dura (Paxinos and Watson, 1998). Acrylic dental cement adhered to skull screws was used to stabilize the assembly on the animal's head. All surgical procedures have been reviewed and approved by the University of Otago Animal Welfare Office.

Postoperative training. After 1 week of recovery, rats were reintroduced to the maze with head plugs connected to a tethered head stage that housed two light-emitting diodes (LEDs) for tracking. In the initial week of postoperative training, four cereal pellets were placed at each reward site, and rats were retrained to run the maze in continuous, unidirectional paths as described above. All 10 animals quickly recalled the task with no indication of postoperative motor impairments. As before, reversals in unidirectional travel or attempts to circumnavigate the midstem were blocked by the experimenter with a particle board insert, but at this stage of training, such corrections were minimally required. In the second week of postoperative training, four cereal pellets were placed at each reward site, and 15 cm barriers were reintroduced to both arms of the maze. Barrier height was increased 5 cm/d until the full height of 30 cm was achieved. During both preoperative and postoperative training periods, when cost–benefit values were equalized between the two arms of the maze (4B:4B configuration), animals distributed their behavior fairly evenly between right and left choices. During the recording sessions, the HCHR arm (barrier location to right or left) was counterbalanced between subjects, although herein the HCHR location has been consistently illustrated on the right for clarity. Each testing session comprised 60 trials; when animals ran for >70 trials, they began to spend long periods whisking and grooming rather than in pursuit of food reward.

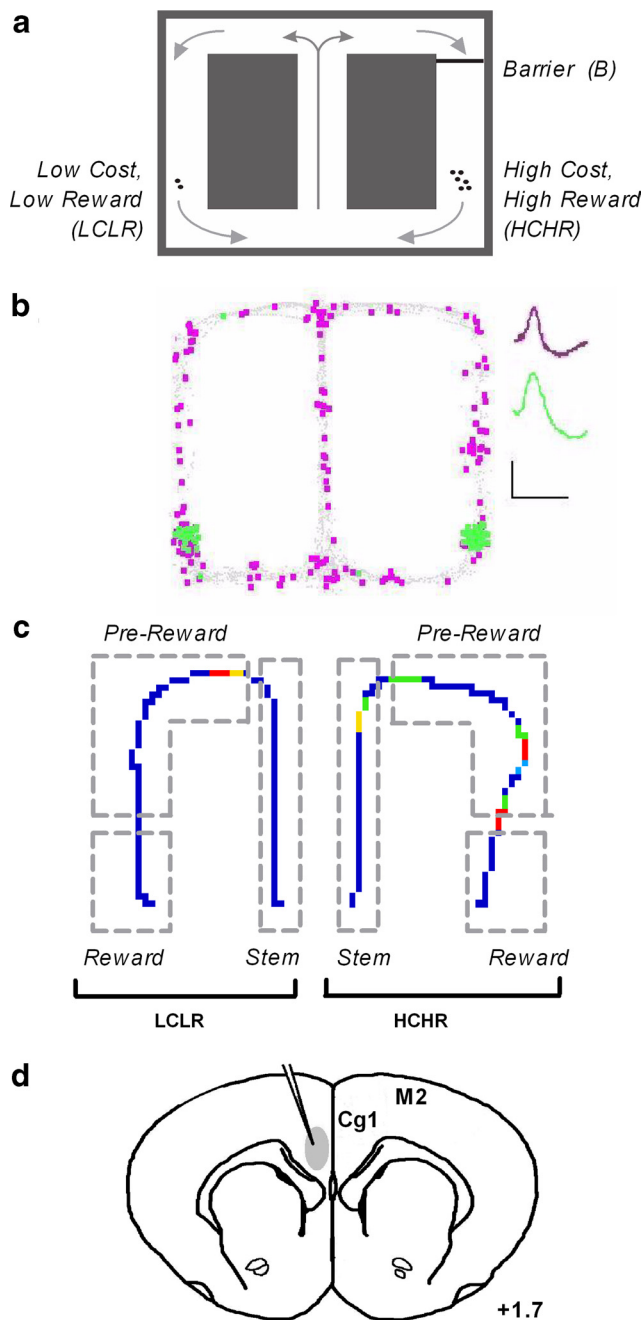


Figure 1. LCLR/HCHR discrimination task. *a*, Diagram of the 2:6B configuration. For each trial, the animal ran up the midstem of the maze and then opted to pursue either two cereal pellets (LCLR) or six cereal pellets (HCHR), the latter obtained only after scaling a 30 cm barrier. After consuming the cereal, the rat continued along the base of the maze and turned at the midstem to initiate another trial. *b*, Distribution of firing in the maze for two example neurons recorded during one animal's session. Most ACC neurons we recorded from (93%) exhibited distributed firing throughout the maze, similar to the cell whose spikes are shown here in magenta. Four neurons fired predominately in the reward region, similar to the cell shown here in green. Mean waveforms are provided on the right. Calibration: 1 ms, 100 μ V. *c*, For comparative analysis of neuronal firing during the task, regions of interest were defined to encompass the stem, prereward, and reward areas for each trajectory. For most analyses, the critical comparison was between trajectories. Two sample trajectories are shown, one to each side, and firing rate pseudo-color mapped onto position (hotter color corresponds to higher firing rate). *d*, Schematic coronal section of the rat brain at +1.7 mm from bregma. Histological analysis confirmed that all recording electrodes were within the ACC; the shaded oval encompasses the electrode tip locations from all 10 animals. Cg1, Cingulate; M2, secondary motor cortex.

Data acquisition. Neuronal activity was acquired and processed using the DacqUSB data acquisition system (Axona). Signals were digitized with 48 kHz sampling, amplified, and bandpass filtered at 360–7000 Hz. Animal movement was tracked and recorded via integrated LED tracking sampled at 10 Hz. Data were only included for analysis if the animal sampled both choice arms throughout the session and navigated in a goal-directed manner without extended grooming, rearing, or immobility. Eight sessions from one rat were excluded because the animal was averse to climbing the barrier; these excluded sessions involved long periods of immobility in front of the barrier or 100% low-cost, low-reward (LCLR) selection. For offline analysis, spikes were sorted and clustered into cell-specific groupings using Tint software (Axona). Cluster and waveform characteristics were used to identify the same cell across multiple days of recording. Firing rates were extracted with in-house MapInfo software using the regions of interest outlined in Figure 1c. For each pass around the maze, three firing rate values were generated: one value for the stem, one value for the prereward region, and one value for the reward region. Firing rates were normalized for region. The mean firing rate across all cells was 1.69 ± 0.16 Hz, which is consistent with previous characterizations of principle cells in the ACC (Jung et al., 1998; Gemmell et al., 2002; Fujisawa et al., 2008). Regional data were exported to a spreadsheet and analyzed using Prism 5.0 (GraphPad Software). Mean firing rates for each region of interest were calculated using composite data from all 60 trials. Regional mean firing rates for the entire session were often < 2 Hz as a result of some trials containing minimal to no ACC activity, particularly those on which the rat repeatedly selected the same side. Regional differences in firing were assessed using a repeated-measure ANOVA, with factors of region (stem, prereward, or reward) and trajectory (LCLR or HCHR). For experiments involving a configuration manipulation, a third factor of configuration was included. Firing rate biases were calculated using the average firing for each trajectory (stem + prereward + reward) over a block of n trials.

Histology. After completion of the study, animals were deeply anesthetized with halothane, and recording sites marked with direct current (2 mA for 5 s) before transcardial perfusion. After fixation, prepared coronal sections were stained with thionin and visually inspected to determine electrode placement.

Results

Ten male Sprague Dawley rats were trained in a cost–benefit, continuous T-maze task that has been shown previously to be sensitive to ACC lesions (Walton et al., 2003; Schweimer et al., 2005). In the baseline condition, animals chose between one arm of the maze that contained two cereal pellets (an LCLR option) and one arm that contained six cereal pellets positioned behind a 30 cm scalable barrier (an HCHR option) (Fig. 1a). We will refer to this configuration as “2:6B” to denote pellet ratio and barrier location. Rats were allowed 60 free choice trials per test phase during which all animals sampled both choice arms. Mean \pm SEM running speed over the whole maze was 14.6 ± 3.8 cm/s (range of 9.1–21 cm/s). Midstem running speed was no different ($p > 0.22$) between LCLR and HCHR choices for any of the 10 animals. From 124 recording sessions, 54 neurons were stably recorded for the full duration of the multiday experiments. The majority of these neurons ($n = 50$) exhibited firing throughout all parts of the maze, although four cells fired almost exclusively at the reward locations (Fig. 1b). Across all cells, the mean \pm SEM firing rate over the whole maze was 1.69 ± 0.16 Hz (range of 0.10–8.9 Hz); the mean \pm SEM maximum firing rate was 8.7 ± 0.78 Hz (range of 3.1–16.8 Hz). These firing rates are consistent with previous reports for pyramidal neurons in the rat ACC (Jung et al., 1998; Gemmell et al., 2002; Fujisawa et al., 2008). There was no indication that the firing rates of the four “reward” cells were different from the population as a whole. To determine whether the firing rates of individual cells varied depending on where the animal was in the maze, each trial trajectory was divided into

three regions of interest matched between LCLR and HCHR sides to make six regions in total (Fig. 1c). Larger regions of interest were used given the spatial variability in firing. Histological analysis confirmed that all recording electrodes were located in the ACC (Fig. 1d).

ACC neurons exhibit biased firing favoring the HCHR arm

In our baseline LCLR/HCHR discrimination task (2:6B), seven animals preferred the HCHR option (HCHR choices: mean \pm SEM of $73 \pm 1.9\%$, range of 71–82%, Wilcoxon’s signed rank test, $p \leq 0.03$), whereas three animals lacked a significant preference ($61.8 \pm 3.4\%$, 47–65%, $p \geq 0.17$). For all subsequent analyses, animals were treated as one group. When firing rates were analyzed across the six regions of interest outlined in Figure 1c, there was not a consistent significant main effect of region (stem, prereward, or reward) within any trajectory. Some cells exhibited localized firing within the stem and prebarrier areas, but there was a high degree of cell-to-cell variability that precluded a definitive localization pattern. Despite this variability in regional firing (stem, prereward, or reward), a substantial number of ACC neurons ($n = 32$ of 54) displayed significantly higher firing in HCHR-associated regions compared with matched LCLR regions (all $F_{(1,147)} \geq 31.8$, $p < 0.01$, main effect of trajectory) (Fig. 2). Firing rate trajectory biases were not attributable to differences in running speed, turn direction, or barrier location, and, interestingly, HCHR-biased cells were found in all 10 animals, even those that lacked an overt behavioral preference for the HCHR option. There was no indication that HCHR-preferring animals contributed more cells toward the cell total than nonpreferring animals. This finding of a trajectory bias suggests that these neurons may encode some feature of the current path beyond simple turn preference. Of the remaining 22 cells that were not HCHR biased, 20 cells lacked a significant effect of trajectory and only two were LCLR biased.

The increased firing rate observed in many of these cells on HCHR trials compared with LCLR trials could be attributed to several factors. First, there is greater physical effort required to gain reward in the HCHR arm of the maze. Anticipation and/or execution of the motor skills required to climb the barrier may account for the increase in ACC activity. Increased ACC activity on HCHR choices could also reflect anticipation and/or consumption of a larger quantity of food on this side. Although there was no clear indication that the firing rate bias occurred preferentially in one or more of the three regions of interest, to determine whether effort (barrier climbing) or reward (food quantity) directly influenced ACC firing, we independently manipulated these variables over 3 consecutive days of recordings (Fig. 3). Sixteen cells were recorded in total from five animals. Twelve of these cells were preselected as responsive cells given that they exhibited significant differential firing (all $F_{(1,147)} \geq 24.8$, $p < 0.01$, main effect of trajectory) for the HCHR trajectory when tested in a 2:6B screening condition. The remaining four cells lacked a trajectory bias in the 2:6B condition, and their firing rates did not significantly change in response to any of the manipulations presented over the 3 d (data not shown).

On day 1, all 12 responsive cells were initially tested for 25 trials in a 2:6 condition in which the choice was solely between differences in reward magnitude. Surprisingly, despite the fact that all of these cells had shown differential firing in the 2:6B prescreening condition, only one cell displayed differential firing in the 2:6 configuration. All five animals appeared to detect the difference in reward value as evidenced by their behavioral preference (mean \pm SEM of HR choices, $84 \pm 3.6\%$, range of 72–

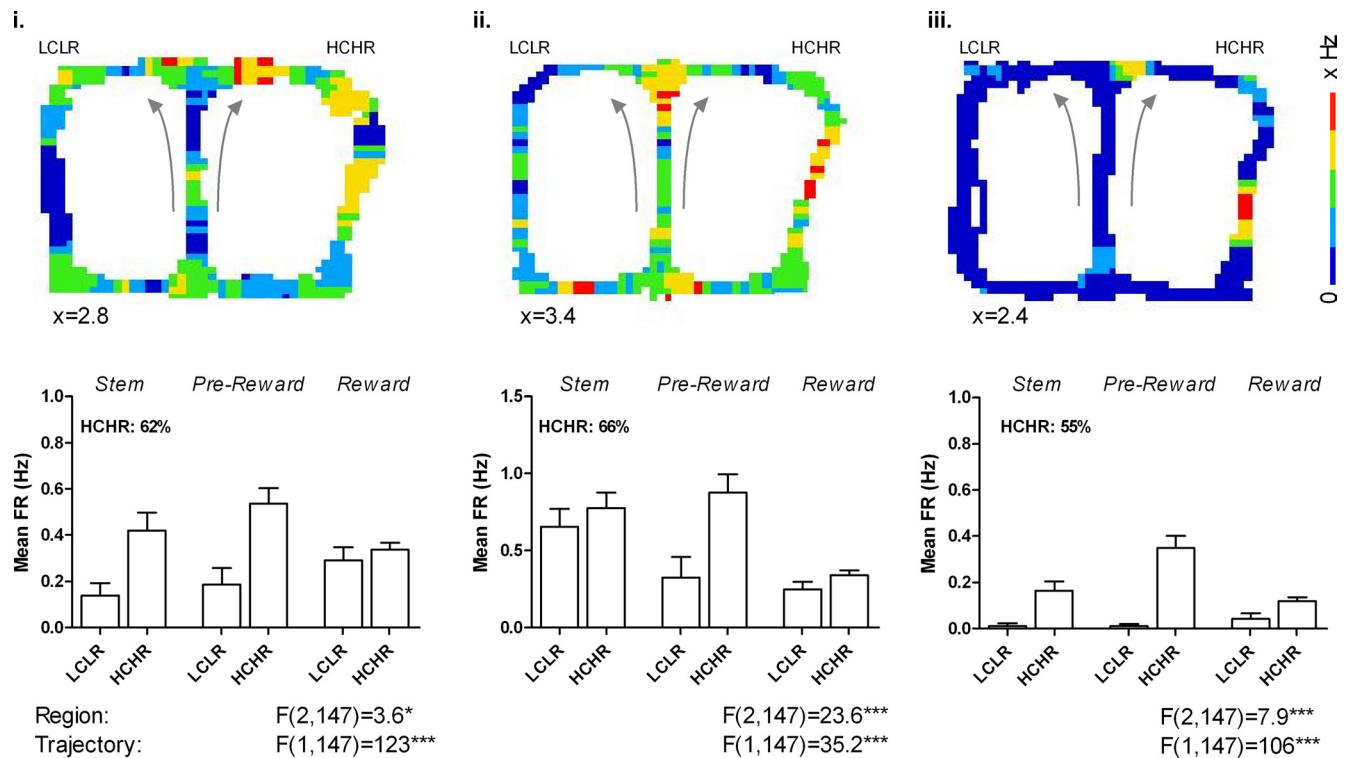


Figure 2. Main effect of trajectory on ACC firing. A substantial portion of ACC neurons (63%) had significantly higher firing throughout one trajectory compared with the other; for 94% of these cells, higher firing was associated with the HCHR option. Data from three typical HCHR-biased cells (*i–iii*) are shown here. For each cell, firing rate data across a 60 trial session are pseudo-color mapped onto position in the maze to visually depict firing rate differences between the trajectories. Firing rate scale values are given under each color map. Note that for some LCLR traversals, particularly in cell *iii*, ACC firing is minimal. Composite data (mean firing rate \pm SEM) from all 60 trials is shown graphically, and illustrates the overall bias toward higher firing in HCHR-associated regions. Behavior is given as percentage of HCHR choices. The ANOVA results for region and trajectory factors are shown under each graph (* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$ here and in subsequent figures).

92%). The responsive cell was one of the minority of cells that fired predominately in the reward zones, and, in the 2:6 condition, it exhibited significantly higher firing in the HR compared with the LR region. The remaining 11 cells had low firing rates overall with no difference between LR and HR trajectories (composite data in Fig. 3*a*, left). When a barrier was inserted into the HR arm of the maze on trial 26, to reinstate a 2:6B condition, firing rates for all 12 cells increased and were again significantly higher in the HCHR compared with LCLR trajectory (Fig. 3*b*, left).

On day 2, the 2:6 condition was repeated for the first 25 trials, and the results of day 1 were replicated (Fig. 3*a*, middle). When barriers were inserted into both sides of the maze on trial 26 to create a 2B:6B condition, the firing rates for all 12 cells were significantly higher ($p < 0.01$, main effect of trajectory) for choices to the 6B side, indicating that it was not the barrier per se that was driving the increased firing rate (Fig. 3*b*, middle). On day 3, the initial condition was 2:2 for the first 25 trials, and there was no differential firing evident in any cells (Fig. 3*a*, right). When a 2:2B condition was introduced on trial 26 (Fig. 3*b*, right), mean firing rates of all cells were now significantly higher ($p < 0.01$, main effect of trajectory) on trials through the nonbarrier arm of the maze. Together, these results suggest that ACC activity does not directly reflect the cost (physical effort) or benefit (food quantity) of an available course of action but rather appears to represent a combinatorial cost–benefit computation biased toward what may be the economically advantageous side given the current context. Furthermore, these data suggest that some threshold level of effort is required before ACC neurons become responsive to these factors because, when little effort was required

to solve the task (2:2), ACC activity was minimal, even when a reward differential existed (2:6).

ACC encoding of cost–benefit is relative and dynamic

Optimal decision-making requires a relative comparison of available options. In our study, for example, the two-pellet (no-barrier) option could be considered the inferior choice in the 2:6 configuration but the superior choice in the 2:2B configuration. Thus, for a neural representation of cost–benefit to be of optimal use to an animal, the encoding should respond to currently available options in a relative and flexible manner. The data presented in Figure 3 indicate that this is the case for ACC neurons; firing rates adapted to changes in cost–benefit ratios, seemingly favoring what may be the “better” of the two options, and did not simply respond directly to reward magnitude or barrier location in isolation. It is possible, however, that the sudden insertion of barriers into the maze during the shift from the baseline to the barrier configuration may have triggered a response to novelty rather than a cost–benefit calculation. To test this possibility, we designed two additional dynamic cost–benefit tasks in which the barrier placement remained unchanged: (1) a multiday experiment in which we presented intersession changes in cost/benefit ratios, and (2) a single day experiment in which we presented intrasession changes in cost/benefit ratios. In both experiments, the 2:6B configuration served as the baseline condition.

In the first experiment, rats were given 60 trials/d in stable LCLR/HCHR discrimination tasks; configurations are outlined at the top of Figure 4; an example of the firing pattern of one cell is illustrated in Figure 4*a*. On day 1 when the configuration was

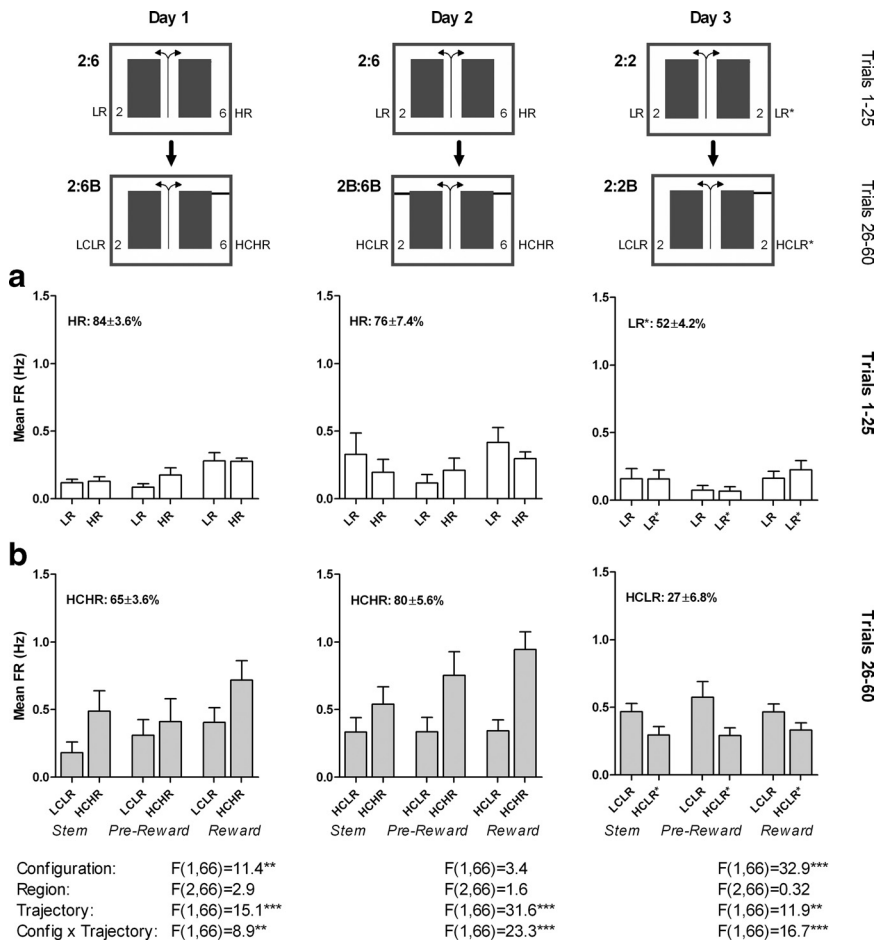


Figure 3. Independent manipulations of cost and benefit. On each day of this 3 d experiment, rats were given 25 trials in a baseline, no-barrier condition (white bars), followed immediately by 35 trials involving one or two barriers (gray bars). Configurations for each day are outlined at the top of the figure, and the corresponding data are presented below. Twelve ACC neurons from five animals were stably recorded for the full course of the experiment. Composite data are given as mean \pm SEM firing rate. Behavioral data are given as mean \pm SEM percentage of HR choices. ANOVA results are presented at the bottom. **a**, In the baseline, no-barrier conditions, mean ACC firing rates were minimal and no different between the two trajectories of the maze ($p > 0.1$, ANOVA), even when a difference in reward magnitude was present (days 1 and 2). **b**, When barriers were inserted into the maze, a significant main effect of trajectory and an interaction between configuration and trajectory were observed. For days 1 and 2, firing rates were biased toward the 6B side after the manipulation. For day 3, the manipulation produced a bias toward the two, nonbarrier side. Note that, across the 3 d experiment, high ACC firing did not consistently correlate to the high-reward (6 pellet) arm nor the high-cost (barrier) arm but did correspond to what might be interpreted as the better side given the cost–benefit configuration.

2:6B, the majority of recorded neurons ($n = 10$ of 14, 4 animals) exhibited significant differential firing ($p < 0.01$, main effect of trajectory) favoring the HCHR trajectory. Composite data from these 10 responsive cells are shown in Figure 4b. The remaining four neurons showed no bias to either side (data not shown). On day 2 when the reward configuration changed to 2:2B, all 10 responsive neurons again exhibited significant differential firing for the trajectory factor, but the higher firing rate now favored the LCLR side in all cases. This suggests that individual ACC neurons have a capacity to flexibly and relatively encode the different cost–benefit configuration. On day 3 when the reward configuration was a spatially reversed 2:6B, all 10 neurons shifted their firing rate bias toward the HCHR side. This trajectory bias developed despite the animal being challenged with the reversal, whereby the turn direction to the HCHR arm was opposite to that of day 1. To determine the time course required to establish this relative cost–benefit encoding, we calculated the firing rate ratio for each cell (FR_{HC}/FR_{LC}) and plotted this against the behavioral

choice of the animal (percentage of HC choices) for every 10 trials throughout each session. Mean data for all 10 cells is shown in Figure 4c. When analyzed as a population, there was a significant correlation between firing rate bias and behavior for each day of the protocol. Firing rate biases favoring the economically advantageous side in the configuration for that day appeared to develop progressively and gradually over the recording session, paralleling the change in choice behavior. It should be noted that, when analyzed on a cell-by-cell basis, there was not always a significant correlation between firing rate bias and behavioral choice for each configuration. For example, in the 2:6B baseline configuration, 3 of the 10 cells exhibited a significant ($p < 0.05$) R value when analyzed individually. In the 2:2B configuration, in which the effort/reward ratio was altered from baseline, all 10 of the cells exhibited a significant R value. In the 2:6B reversal configuration, in which the turn direction and maze orientation were the opposite of baseline, 7 of the 10 cells exhibited a significant R value.

In the second experiment (Fig. 5), we examined whether ACC neurons could rapidly integrate sudden changes in cost–benefit contingencies that occurred within a recording session. Eighteen neurons were recorded in total from six animals. Twelve were preselected as responsive cells because they exhibited differential firing ($p < 0.01$, main effect of trajectory) in the 2:6B condition. Rats were given 60 trials in a dynamic LCLR/HCHR discrimination task involving a reward reduction (2:6B to 2:2B) on trials 21–40. During the 20 trial reward reduction period, 8 of the 12 responsive neurons demonstrated significant changes in trajectory bias, and the other four showed a trend in the same direction; higher firing rates shifted from HCHR- to LCLR-associated

regions. An example is shown in Figure 5a, and mean data from all 12 cells are shown in Figure 5c. The data demonstrate that changes in ACC encoding can develop rapidly. It is interesting to note that, in this case, the change in firing rate occurred without an overt change in choice behavior (data from one cell, plotted against behavior, are shown in Fig. 5b), indicating that changes in firing rate are not merely a downstream response to behavioral change. Together, these data strongly suggest that ACC neurons encode relative cost–benefit differentials based on the options that are currently available and that this encoding can dynamically adjust in the face of changing circumstances without, or possibly before, a behavioral change.

Discussion

Here we demonstrate that neurons in the rat ACC encode a cost–benefit differential of two competing courses of action, with higher firing rates consistently biased toward the economically advantageous option. Although recent studies in macaque have demonstrated that ACC neurons respond to independent manip-

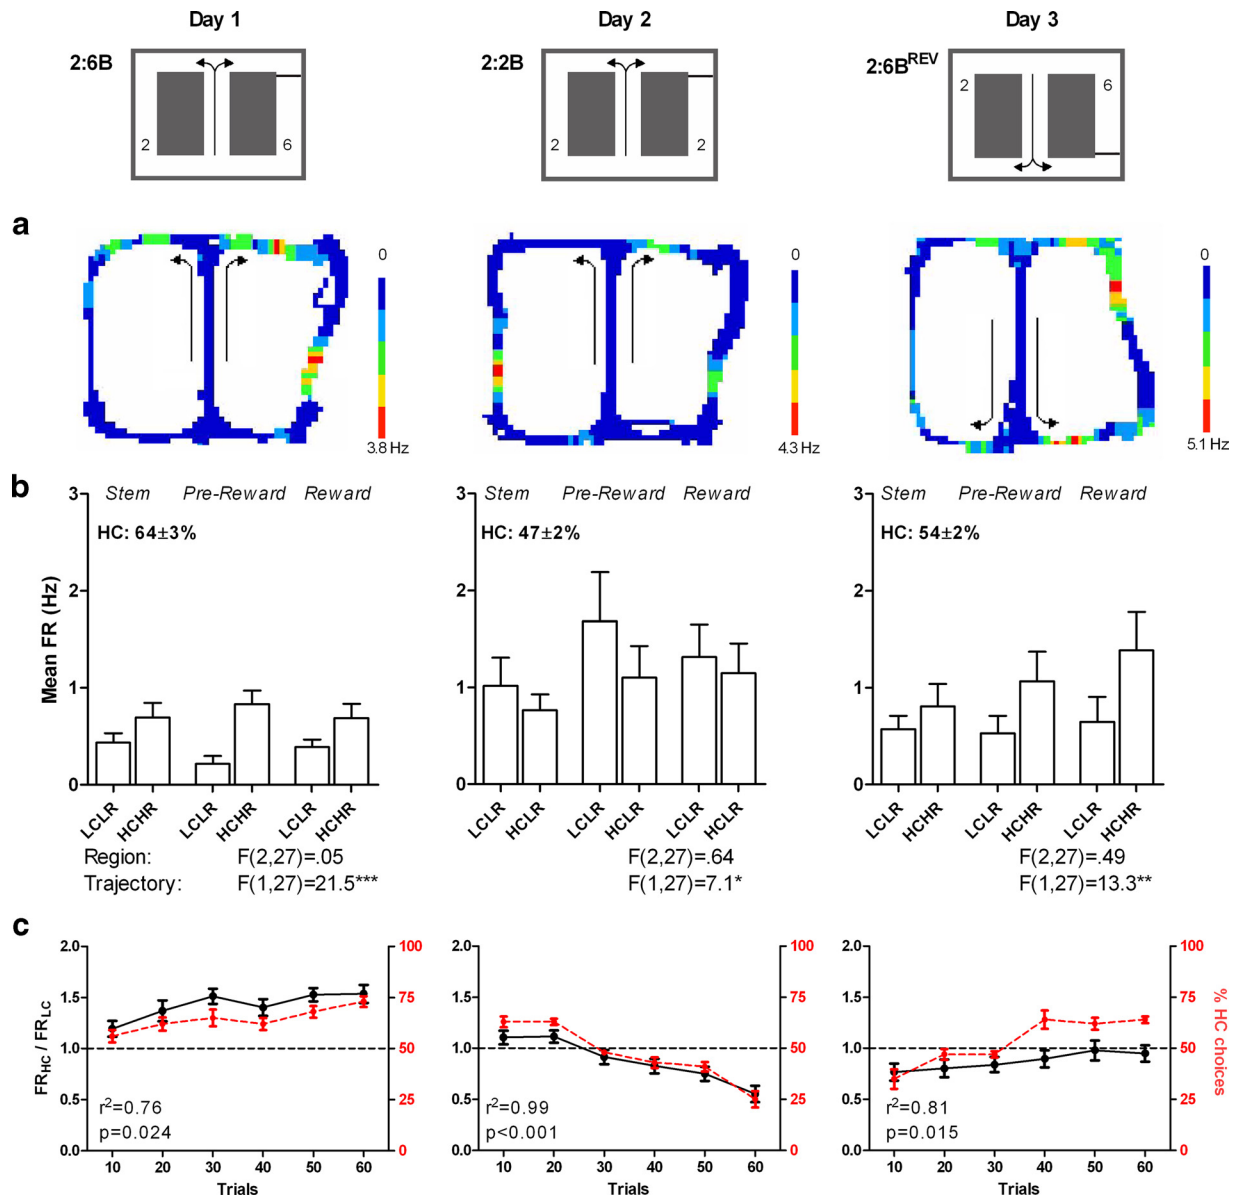


Figure 4. Cost–benefit representations reflect relative encoding. On each day of this 3 d experiment, rats were given 60 trials in the configurations outlined at the top of the figure. Ten ACC neurons from four animals were stably recorded for the full course of the experiment. *a*, Data from one example neuron showing an illustration of firing density for the session of each day. Note that, for days 1 and 3, areas of intense firing are observed at similar points of the HCHR trajectory despite the spatial configuration being reversed. On day 2, intense firing is observed in the LCLR reward zone, despite the same reward (two pellets) being available there on all 3 d. *b*, For all 10 cells, firing rates (mean \pm SEM) were consistently higher for the economically advantageous trajectory, even when animals were challenged with a reward devaluation (day 2) and spatial reversal (day 3). Behavioral data are given as mean \pm SEM percentage of HC choices. ANOVA results are presented below the graph. *c*, For each cell, the firing rate bias between HC and LC trials was computed across every 10 trials for each session. Composite data for all 10 cells are shown here as mean \pm SEM firing rate ratio (black) plotted alongside corresponding mean \pm SEM behavioral choice (red). Note that firing ratios become biased toward the economically advantageous side relative to the configuration of the day as the session progresses. R^2 and p values indicate the fit between mean firing rate bias and behavior.

ulations in reward size or effort cost (Kennerley and Wallis, 2009a; Kennerley et al., 2009), this is the first study to show that ACC activity is strongly modulated in situations in which an integration of effort and reward is needed to perform optimally. Our data suggest that, when the physical demands of a task are minimal, ACC activity is low and nondiscriminatory, even in cases in which there is a clear difference in reward magnitude. Once a moderate level of physical effort is required to achieve one or more goals, however, ACC activity increases and firing rate trajectory biases develop. Separate choice options do not appear to be represented by separate cell subpopulations; rather, the same population differentially encodes each trajectory with variations in firing rate. Our findings suggest that the key role of the

ACC may be to provide information about choice value when some level of effort must be expended in the process of realizing that choice.

The present results indicate that ACC neurons encode cost–benefit in relative and dynamic terms, two characteristics that confer flexibility to the decision-making process. ACC neurons appear to integrate cost–benefit information using an ordinal scale of utility rather than reflecting absolute reward value or effort cost. For example, firing rate biases associated with the two-pellet, nonbarrier option were different depending on whether that option was presented in a 2:2B versus a 2:6B condition (Figs. 3, 4). An encoding system that operates on relative value is essential for identifying the best choice among available

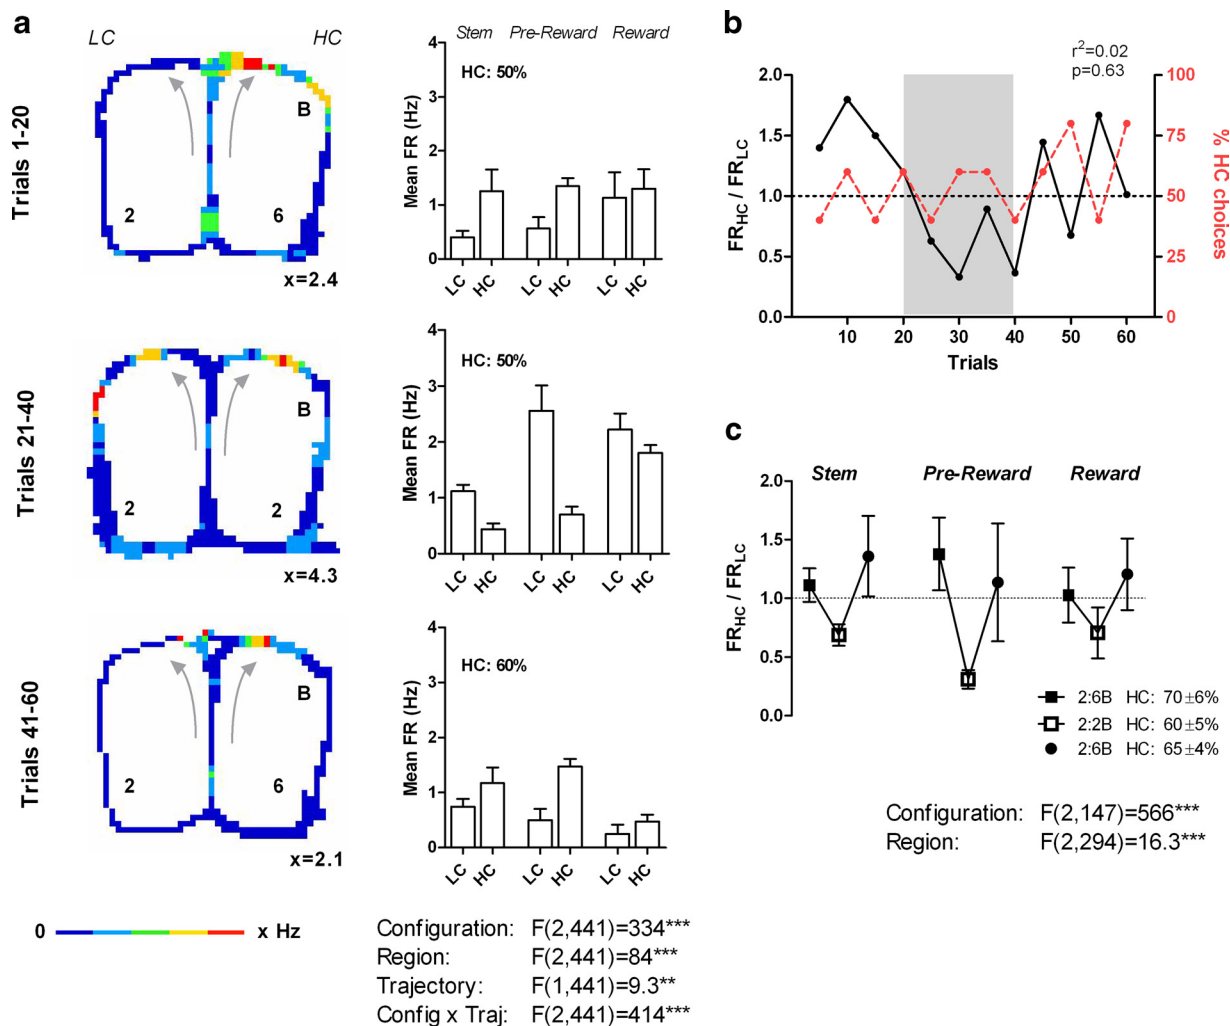


Figure 5. ACC encoding rapidly adapts to changes in the environment. In this dynamic task, rats were given 60 trials in which a reward reduction occurred on trials 21–40. Twelve neurons from six animals were recorded. **a**, Results from one individual cell illustrating mean firing rate for trials 1–20 (top), 21–40 (middle), and 41–60 (bottom). A robust shift in firing rate bias from the HC to the LC side occurred during the reward reduction phase. ANOVA results for this neuron are presented below the last graph. **b**, Time course of the data from the neuron shown in **a** depicting the shift in firing rate bias during the session. Firing rate ratios for this individual cell were calculated for every five trials (black) and plotted against behavior (red). During the 2:2B phase (shaded), there was a significant decrease in firing rate ratio compared with the initial 2:6B phase ($F_{(2,9)} = 8.1, p = 0.0097$). No correlation between behavior and firing rate bias was observed for this cell. **c**, Composite data from all 12 neurons, illustrating the firing rate bias during the 2:6B (■), 2:2B (□), and second 2:6B (●) configuration for each defined region of interest. Mean \pm SEM behavioral data are provided as percentage of HC choices. Note the main effect of configuration on firing rate bias.

alternatives and moreover enables efficient cortical computation and an ability to deal with novelty. The relative encoding we observed is similar to that reported in monkey ACC during a scaled-reward task (Sallet et al., 2007), and an analogous notion of an abstract value system has been proposed by others (Amiez et al., 2006; Wallis, 2007; Kennerley et al., 2009). The finding that relative encoding can be readily elicited and observed in rats in this task suggests that this may be a useful new model system for studying how choice value is represented in the brain.

ACC neurons also appear to integrate cost–benefit information in a dynamic manner, as suggested by the trajectory bias shifts we observed in response to intersession and intrasession changes in cost–benefit configuration (Figs. 3–5). Lesion studies in monkey report that the ACC is needed to flexibly adapt choice behavior to maximize reward gain (Shima and Tanji, 1998; Rushworth et al., 2003; Amiez et al., 2006), and human imaging shows an increase in ACC activity when task reward values diminish and action shifts are subsequently more likely to occur (Bush et al., 2002). Previous single-unit studies have shown that, on a trial-by-trial

basis, ACC and cingulate motor neurons respond to reward receipt, reward omission, and changes in expected reward magnitude (Shima and Tanji, 1998; Sallet et al., 2007; Quilodran et al., 2008). Here we show that, if an animal routinely samples choice options and a change in reward or effort is experienced, ACC neurons can promptly shift their firing rate bias in subsequent trials toward what may be interpreted as the “better” option given the current cost–benefit configuration.

Although we found a consistent pattern of higher firing corresponding to the better side of each barrier configuration, animal behavior did not always directly mirror ACC firing rate biases on a trial-by-trial basis. For instance, in the intrasession manipulation presented in Figure 5, the firing rate bias shifted toward the economically advantageous low-effort option during the reward reduction period with no change in choice behavior. In the intersession manipulations presented in Figure 4, overall firing rate biases from each session favored the advantageous option, but in the novel 2:2B and 2:6B^{rev} configurations, more cells exhibited significant firing rate bias–behavior correlations

than in the 2:6B baseline configuration. If firing rate biases had consistently mimicked behavior, it could be argued that ACC activity was simply a corollary of motor output. Rather, our data suggest that the ACC is continually updating a cost–benefit representation of the task set that may only affect behavior when increased attention to choice/adaptation is warranted. Imaging studies are consistent in suggesting that the ACC frequently updates task sets even if available options are not used immediately (Hyafil et al., 2009). In Figure 5, behavioral adaptation in response to the reward reduction would not necessarily be expected, or optimal, given that an element of volatility has been introduced to the task. It is possible that the ACC is continually responsive to cost–benefit changes but that choice behavior is dependent on population coding, with individual neurons responding to cost–benefit changes with different levels of sensitivity and/or different integration times. Only when sufficient neurons have switched their cost–benefit “preference” would behavior be affected. Population coding aligns with the findings of temporal reward signal filtering demonstrated in monkey (Seo and Lee, 2007), and network state dynamics within the ACC appear to underlie successful decision-making in a radial maze task (Lapish et al., 2008).

In the baseline 2:6B configuration of our task, our animals never completely shifted their behavior toward the HCHR option so as to maximally exploit food reward. This is in contrast to many previous ACC studies in which animals chose the HCHR option on 90% or more of trials in the baseline configuration (Walton et al., 2002; Schweimer et al., 2005; Rudebeck et al., 2006), likely shaped by preoperative performance criteria. One interpretation of our animals’ behavior is that, with intermediate levels of training and no forced choices, the rats were more likely to incorporate routine sampling into their behavior in attempts to monitor choice options throughout the study. Three of our rats routinely alternated choice behavior at the onset of each session, perhaps as an exploratory mechanism to establish the cost–benefit configuration of the day. As is increasingly appreciated in the decision-making literature, sampling is a critical component of optimal decision-making in a dynamic world. Gambling studies, for example, demonstrate that human subjects will forfeit one choice trial on a known high-reward option in an effort to gain information on the alternative option (Daw et al., 2006).

Compared with the results of other studies (Amiez et al., 2006; Sallet et al., 2007; Seo and Lee, 2007; Kennerley et al., 2009), we found little evidence of differential encoding in our 2:6 discrimination task despite the difference in reward magnitude and the animal’s behavioral preference (Fig. 3*a*). Rather, our findings indicate that there is a modulating influence of effort that determines whether or not reward-related activity is observed in the ACC. This suggests that, in previous primate studies, the animals likely registered a component of effort associated with the task. This could occur even if physical effort requirements were equal between action sets but increased mental effort was needed to perform optimally. For example, when rewarded actions must be determined by trial and error (Procyk et al., 2000; Quilodran et al., 2008), the process demands heightened cognitive effort (e.g., working memory load). In such tasks, ACC neurons are found to differentially encode the expected value of each choice, but activity decreases once the optimal choice is identified and repeated (Procyk et al., 2000). Cognitive effort may therefore modulate ACC encoding in primates similar to how physical effort modulated encoding in our rodent study. Imaging in humans supports this notion of ACC recruitment by cognitive effort; higher ACC

activity is observed in tasks requiring higher cognitive effort, such as those presenting conflict (Botvinick, 2007), volatility (Behrens et al., 2007), or distracters (Barch et al., 1997; Fu et al., 2002).

In conjunction with previous reports, our data suggest that choice behavior based solely on established reward information requires minimal ACC activity, which reflects the ability of ACC-lesioned animals to still perform at control levels in a variety of decision-making tasks (Amiez et al., 2006; Kennerley et al., 2006; Rudebeck et al., 2006). Reward location and path planning, for example, appear to rely more on the prelimbic/infralimbic region of the prefrontal cortex than the ACC (Kesner and Ragozzino, 2003; Hok et al., 2005; Kennerley and Wallis, 2009b). This is an important reminder that the ACC is just one part of a decision-making network incorporating but not limited to the lateral prefrontal cortex, orbitofrontal cortex, basolateral amygdala, and striatum (Salamone, 1994; Schultz, 2000; Floresco and Ghods-Sharifi, 2007; Croxson et al., 2009; Hauber and Sommer, 2009). Although a unifying perspective of the role of the ACC in this decision-making network has not been agreed on, we propose that the ACC is specifically recruited for cost–benefit integration when a decision task requires a threshold level of mental or physical effort to optimally perform. This may involve a serial transfer of information between the nucleus accumbens and the ACC (Hauber and Sommer, 2009), with ACC action–outcome information energizing the lateral prefrontal cortex to guide task-appropriate behavior (Kouneiher et al., 2009).

References

- Amiez C, Joseph JP, Procyk E (2006) Reward encoding in the monkey anterior cingulate cortex. *Cereb Cortex* 16:1040–1055.
- Barch DM, Braver TS, Nystrom LE, Forman SD, Noll DC, Cohen JD (1997) Dissociating working memory from task difficulty in human prefrontal cortex. *Neuropsychologia* 35:1373–1380.
- Behrens TE, Woolrich MW, Walton ME, Rushworth MF (2007) Learning the value of information in an uncertain world. *Nat Neurosci* 10:1214–1221.
- Bilkey DK, Russell N, Colombo M (2003) A lightweight microdrive for single-unit recording in freely moving rats and pigeons. *Methods* 30:152–158.
- Botvinick MM (2007) Conflict monitoring and decision making: reconciling two perspectives on anterior cingulate function. *Cogn Affect Behav Neurosci* 7:356–366.
- Bush G, Vogt BA, Holmes J, Dale AM, Greve D, Jenike MA, Rosen BR (2002) Dorsal anterior cingulate cortex: a role in reward-based decision making. *Proc Natl Acad Sci U S A* 99:523–528.
- Carmichael ST, Price JL (1995a) Limbic connections of the orbital and medial prefrontal cortex in macaque monkeys. *J Comp Neurol* 363:615–641.
- Carmichael ST, Price JL (1995b) Sensory and premotor connections of the orbital and medial prefrontal cortex of macaque monkeys. *J Comp Neurol* 363:642–664.
- Croxson PL, Walton ME, O’Reilly JX, Behrens TE, Rushworth MF (2009) Effort-based cost–benefit valuation and the human brain. *J Neurosci* 29:4531–4541.
- Daw ND, O’Doherty JP, Dayan P, Seymour B, Dolan RJ (2006) Cortical substrates for exploratory decisions in humans. *Nature* 441:876–879.
- Floresco SB, Ghods-Sharifi S (2007) Amygdala-prefrontal cortical circuitry regulates effort-based decision making. *Cereb Cortex* 17:251–260.
- Forstmann BU, Brass M, Koch I, von Cramon DY (2006) Voluntary selection of task sets revealed by functional magnetic resonance imaging. *J Cogn Neurosci* 18:388–398.
- Fu CH, Morgan K, Suckling J, Williams SC, Andrew C, Vythelingum GN, McGuire PK (2002) A functional magnetic resonance imaging study of overt letter verbal fluency using a clustered acquisition sequence: greater anterior cingulate activation with increased task demand. *Neuroimage* 17:871–879.
- Fujisawa S, Amarasingham A, Harrison MT, Buzsáki G (2008) Behavior-dependent short-term assembly dynamics in the medial prefrontal cortex. *Nat Neurosci* 11:823–833.

- Gemmel C, Anderson M, O'Mara SM (2002) Deep layer prefrontal cortex unit discharge in a cue-controlled open-field environment in the freely-moving rat. *Behav Brain Res* 133:1–10.
- Hauber W, Sommer S (2009) Prefrontostriatal circuitry regulates effort-related decision making. *Cereb Cortex* 19:2240–2247.
- Hok V, Save E, Lenck-Santini PP, Poucet B (2005) Coding for spatial goals in the prelimbic/infralimbic area of the rat frontal cortex. *Proc Natl Acad Sci U S A* 102:4602–4607.
- Hoover WB, Vertes RP (2007) Anatomical analysis of afferent projections to the medial prefrontal cortex in the rat. *Brain Struct Funct* 212:149–179.
- Hyafil A, Summerfield C, Koehlin E (2009) Two mechanisms for task switching in the prefrontal cortex. *J Neurosci* 29:5135–5142.
- Jung MW, Qin Y, McNaughton BL, Barnes CA (1998) Firing characteristics of deep layer neurons in prefrontal cortex in rats performing spatial working memory tasks. *Cereb Cortex* 8:437–450.
- Kennerley SW, Walton ME, Behrens TE, Buckley MJ, Rushworth MF (2006) Optimal decision making and the anterior cingulate cortex. *Nat Neurosci* 9:940–947.
- Kennerley SW, Wallis JD (2009a) Evaluating choices by single neurons in the frontal lobe: outcome value encoded across multiple decision variables. *Eur J Neurosci* 29:2061–2073.
- Kennerley SW, Wallis JD (2009b) Encoding of reward and space during a working memory task in the orbitofrontal cortex and anterior cingulate sulcus. *J Neurophysiol* 102:3352–3364.
- Kennerley SW, Dahmubed AF, Lara AH, Wallis JD (2009) Neurons in the frontal lobe encode the value of multiple decision variables. *J Cogn Neurosci* 21:1162–1178.
- Kesner RP, Ragozzino ME (2003) The role of the prefrontal cortex in object-place learning: a test of the attribute specificity model. *Behav Brain Res* 146:159–165.
- Kouneiher F, Charron S, Koehlin E (2009) Motivation and cognitive control in the human prefrontal cortex. *Nat Neurosci* 12:939–945.
- Lapish CC, Durstewitz D, Chandler LJ, Seamans JK (2008) Successful choice behavior is associated with distinct and coherent network states in anterior cingulate cortex. *Proc Natl Acad Sci U S A* 105:11963–11968.
- Mars RB, Coles MG, Grol MJ, Holroyd CB, Nieuwenhuis S, Hulstijn W, Toni I (2005) Neural dynamics of error processing in medial frontal cortex. *Neuroimage* 28:1007–1013.
- Paxinos G, Watson S (1998) *The rat brain in stereotaxic coordinates*, Ed 4. San Diego: Academic.
- Procyk E, Tanaka YL, Joseph JP (2000) Anterior cingulate activity during routine and non-routine sequential behaviors in macaques. *Nat Neurosci* 3:502–508.
- Quilodran R, Rothé M, Procyk E (2008) Behavioral shifts and action valuation in the anterior cingulate cortex. *Neuron* 57:314–325.
- Rudebeck PH, Walton ME, Smyth AN, Bannerman DM, Rushworth MF (2006) Separate neural pathways process different decision costs. *Nat Neurosci* 9:1161–1168.
- Rushworth MF, Hadland KA, Gaffan D, Passingham RE (2003) The effect of cingulate cortex lesions on task switching and working memory. *J Cogn Neurosci* 15:338–353.
- Salamone JD (1994) The involvement of nucleus accumbens dopamine in appetitive and aversive motivation. *Behav Brain Res* 61:117–133.
- Sallet J, Quilodran R, Rothé M, Vezoli J, Joseph JP, Procyk E (2007) Expectations, gains, and losses in the anterior cingulate cortex. *Cogn Affect Behav Neurosci* 7:327–336.
- Schultz W (2000) Multiple reward signals in the brain. *Nat Rev Neurosci* 1:199–207.
- Schweimer J, Saft S, Hauber W (2005) Involvement of catecholamine neurotransmission in the rat anterior cingulate in effort-related decision making. *Behav Neurosci* 119:1687–1692.
- Seo H, Lee D (2007) Temporal filtering of reward signals in the dorsal anterior cingulate cortex during a mixed-strategy game. *J Neurosci* 27:8366–8377.
- Shima K, Tanji J (1998) Role for cingulate motor area cells in voluntary movement selection based on reward. *Science* 282:1335–1338.
- Wallis JD (2007) Neuronal mechanisms in prefrontal cortex underlying adaptive choice behavior. *Ann N Y Acad Sci* 1121:447–460.
- Walton ME, Mars RB (2007) Probing human and monkey anterior cingulate cortex in variable environments. *Cogn Affect Behav Neurosci* 7:413–422.
- Walton ME, Bannerman DM, Rushworth MF (2002) The role of rat medial frontal cortex in effort-based decision making. *J Neurosci* 22:10996–11003.
- Walton ME, Bannerman DM, Alterescu K, Rushworth MF (2003) Functional specialization within medial frontal cortex of the anterior cingulate for evaluating effort-related decisions. *J Neurosci* 23:6475–6479.
- Walton ME, Devlin JT, Rushworth MF (2004) Interactions between decision making and performance monitoring within prefrontal cortex. *Nat Neurosci* 7:1259–1265.
- Yoshida W, Ishii S (2006) Resolution of uncertainty in prefrontal cortex. *Neuron* 50:781–789.