

Neurons in Posterior Cingulate Cortex Signal Exploratory Decisions in a Dynamic Multioption Choice Task

John M. Pearson,^{1,*} Benjamin Y. Hayden,¹
Sridhar Raghavachari,¹ and Michael L. Platt^{1,2}

¹Department of Neurobiology, Duke University School of Medicine and Center for Neuroeconomic Studies

²Center for Cognitive Neuroscience and Department of Evolutionary Anthropology
Duke University, Durham, NC 27710, USA

Summary

In dynamic environments, adaptive behavior requires striking a balance between harvesting currently available rewards (exploitation) and gathering information about alternative options (exploration) [1–4]. Such strategic decisions should incorporate not only recent reward history, but also opportunity costs and environmental statistics. Previous neuroimaging [5–8] and neurophysiological [9–13] studies have implicated orbitofrontal cortex, anterior cingulate cortex, and ventral striatum in distinguishing between bouts of exploration and exploitation. Nonetheless, the neuronal mechanisms that underlie strategy selection remain poorly understood. We hypothesized that posterior cingulate cortex (CGp), an area linking reward processing, attention [14], memory [15, 16], and motor control systems [17], mediates the integration of variables such as reward [18], uncertainty [19], and target location [20] that underlie this dynamic balance. Here we show that CGp neurons distinguish between exploratory and exploitative decisions made by monkeys in a dynamic foraging task. Moreover, firing rates of these neurons predict in graded fashion the strategy most likely to be selected on upcoming trials. This encoding is distinct from switching between targets and is independent of the absolute magnitudes of rewards. These observations implicate CGp in the integration of individual outcomes across decision making and the modification of strategy in dynamic environments.

Results

To probe the neuronal processes mediating the strategic balance of immediate reward and information acquisition, we recorded the activity of single cingulate cortex (CGp) neurons in two rhesus macaques performing a “restless” variant of the four-armed bandit for juice rewards [3, 5] (Figure 1). This variant provides a high level of environmental variability with a behaviorally tractable number of options. On each trial, monkeys chose one of four targets whose payoffs were randomly selected from distributions centered about their values on the previous trial. Once a target was chosen, monkeys in principle had perfect knowledge of its present value (there was no added variance in payouts), though the values of all targets changed each trial. As a result, monkeys had to select an option to learn its current value and integrate this information

with their statistical knowledge of the environment to predict its relative value on upcoming trials.

Both monkeys were highly adept at optimizing reward. They earned 92% and 91%, respectively, of the total reward that would have been earned by an omniscient observer. Nevertheless, despite this high level of performance, a perfectly greedy decision maker, focused on the option with highest immediate value, would have harvested more, though not all, available reward (see [Supplemental Data](#) available online).

More importantly, nothing intrinsic to the task design serves to distinguish exploratory from exploitative decisions. On each trial, both monkeys simply selected among the four available options and received a reward. As a result, individual decisions must be classified as exploratory or exploitative according to a model-based analysis of each monkey’s behavior, with model parameters chosen to maximize the likelihood of observed choices. We report here only results based on our best-fitting Kalman filter model, though results were similar for other models as well (see [Supplemental Data](#)).

We analyzed the firing rates of 83 single neurons in CGp in both monkeys performing the four-armed bandit task (59 from monkey N and 24 from monkey B). We focused on two trial epochs, a 2 s decision epoch (DE; 1 s before trial initiation extending to juice delivery) and a 2 s postreward evaluation epoch (EE; from the offset of juice delivery through the intertrial interval). Analyses based on mean firing rates in each epoch readily identified neurons that discriminated between the two strategies (14%, $n = 12/83$, DE; 16%, $n = 13/83$, EE; $p < 0.05$, Mann-Whitney U test), with 22% of neurons doing so in at least one epoch ($n = 18/83$; $p < 0.025$, Bonferroni-corrected Mann-Whitney U test).

Figure 2A depicts the average firing rate of a single neuron on trials classified as either exploratory or exploitative. Responses on exploit trials were significantly higher in both decision and evaluation epochs (Mann-Whitney U test, $p < 0.01$). In contrast, the neuron whose activity is plotted in Figure 2B was more responsive on exploratory trials in both epochs (Mann-Whitney U test, $p < 0.01$). Although the population as a whole exhibited slightly higher firing on exploratory trials in both epochs (modulation index = 0.0084 [DE], 0.0026 [EE]), the population of cells with significant modulation was mixed (modulation index = 0.046 [DE], –0.011 [EE]), indicating heterogeneity in single-cell responses to the different strategies. Thus, firing rates in CGp distinguish between upcoming exploratory and exploitative decisions in the epoch leading up to selection and continue to reflect that choice during the postreward delay.

We previously reported that responses of CGp neurons predict impending switches from one option to another on the next trial in a simple two-alternative task [11]. Based on that finding, we hypothesized that CGp neurons would also carry predictive information about more general impending choices of strategy. To test this hypothesis, we regressed the probability of exploration on the upcoming decision as a function of observed firing rate for each neuron. Of the 83 neurons in our sample, about 16% showed significant correlations between firing rate during the decision epoch and the probability of exploration on the ensuing choice [$n = 13$, $p < 0.05$,

*Correspondence: pearson@neuro.duke.edu

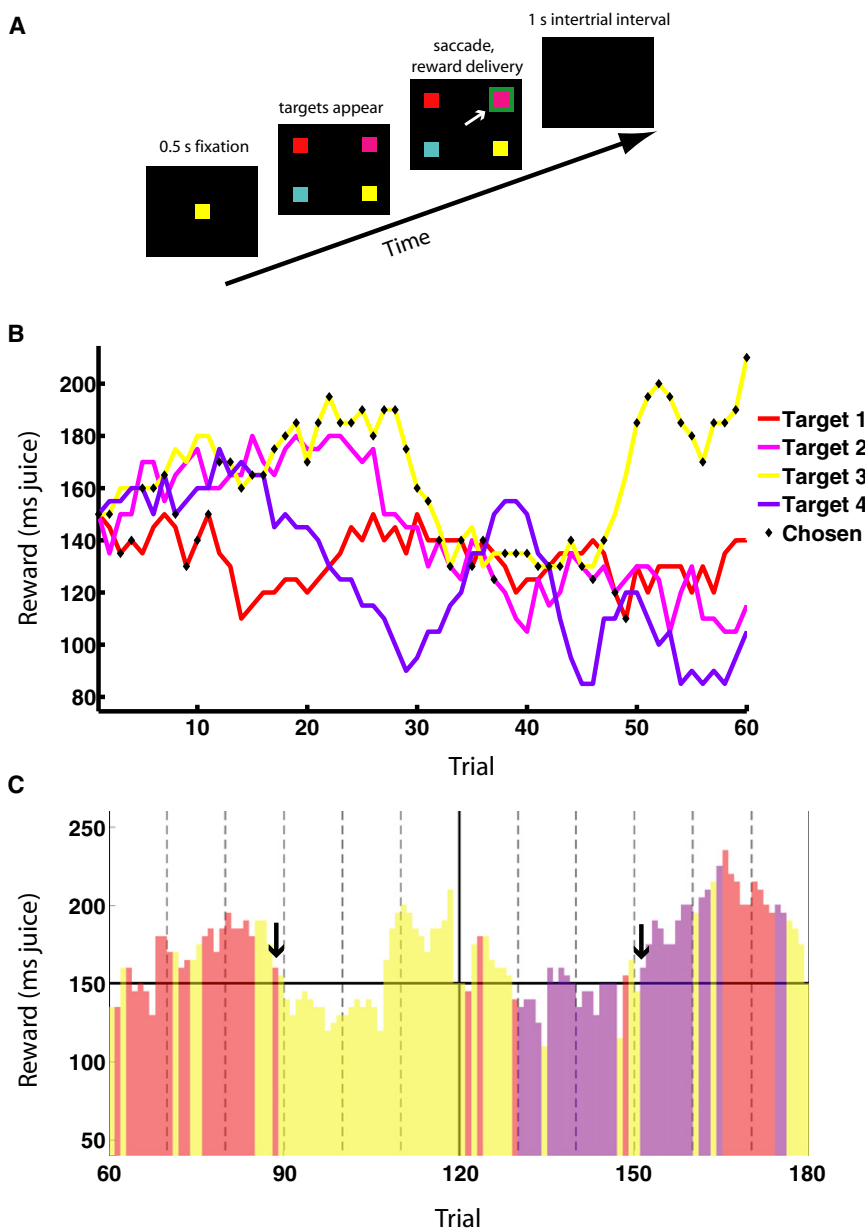


Figure 1. Task and Example Reward Schedule Used to Study the Explore/Exploit Dilemma

(A) Schematic of the four-armed bandit task. Following a 0.5 s fixation period, the central cue disappears, replaced by four colored targets. Subjects indicate choices by shifting gaze to targets, after which the chosen target is highlighted in green for 1 s and a juice reward is delivered. Consecutive trials are separated by a 1 s intertrial interval. Every 60 trials, a block change cue appears, and all target values are reset to the mean reward value.

(B) Sample payouts and choices for the four options over a single block. Reward values for each target follow a random walk with fixed standard deviation for step size, biased toward the mean of 150 ms. Black diamonds indicate choices made by the monkey during the given block.

(C) Sample monkey B choice behavior over two blocks of the four-armed bandit task. Bar colors indicate target chosen; bar heights indicate the values of rewards received. The horizontal line indicates the mean reward value. Monkeys exhibit bouts of exploitation of favorable targets with exploration of alternatives. Arrows indicate trials that might plausibly be classified as either exploratory or exploitative, depending on the behavioral model used. Both involve a change in target selected (action switch) but also a return to a target with high remembered value and so might be classified as exploitative.

to reward value. To investigate this possibility, we calculated and fit reward size tuning curves for each of our 83 neurons (Figures 3A and 3B). Consistent with previous studies [11, 18], we found that the firing rates of CGp neurons during the decision epoch varied with the amount of reward received on the previous trial (DE, $n = 39$, $p < 0.05$, F test for quadratic regression) and that firing rates in the evaluation epoch varied with the amount just received (EE, $n = 44$, $p < 0.05$, F test). Over the range of experienced reward values (50–350 μ l), we found a heterogeneity of

Mann-Whitney U test; $p(n > 12) < 0.001$, binomial test]. Even more importantly, 16% of neurons showed a correlation between firing during EE and the probability of exploration on the following trial [$n = 13$, $p < 0.05$, Mann-Whitney U test; $p(n > 12) < 0.001$, binomial test], suggesting that CGp differentially signals the probability of strategic decisions within a block of trials. Figures 2C and 2D depict the separate population averages for the subsets of cells whose activity correlated negatively ($n = 6$) and positively ($n = 7$) with exploration. Average response for each of these two groups of neurons strongly predicted probability of impending strategic choices in a graded fashion.

One potential confound of these results arises from the link between exploitation and the likelihood of increased reward. Because we might expect that exploitative choices, on average, yielded higher rewards, a possible alternative interpretation of the present data is that effects of strategy on neuronal activity are reducible entirely to neuronal sensitivity

tuning curves: some were linear ($n = 14$ positive, 6 negative, DE; $n = 18$ positive, $n = 9$ negative, EE; $p < 0.05$ nonzero regression coefficient), whereas others were U-shaped, both concave up ($n = 8$, DE; $n = 8$, EE; $p < 0.05$ nonzero regression coefficient) and concave down ($n = 11$, DE; $n = 9$, EE; $p < 0.05$ nonzero regression coefficient). We therefore restricted our next series of analyses to those trials where monkeys adopted different strategies but received the same amount of reward. We found that 12% [$n = 10$; $p(n > 9) < 0.001$, binomial test] of recorded neurons still showed different mean firing rates on explore versus exploit trials ($p < 0.01$, Bonferroni-corrected Mann-Whitney U test). Data for an example neuron showing this effect are shown in Figures 3D–3F. Here, the middle third of reward values have been subdivided into three categories (medium-low, medium-medium, and medium-high), and neuronal firing is plotted as a function of time for both explore and exploit trials, controlled for received reward. This neuron, like many others in our population, showed clear

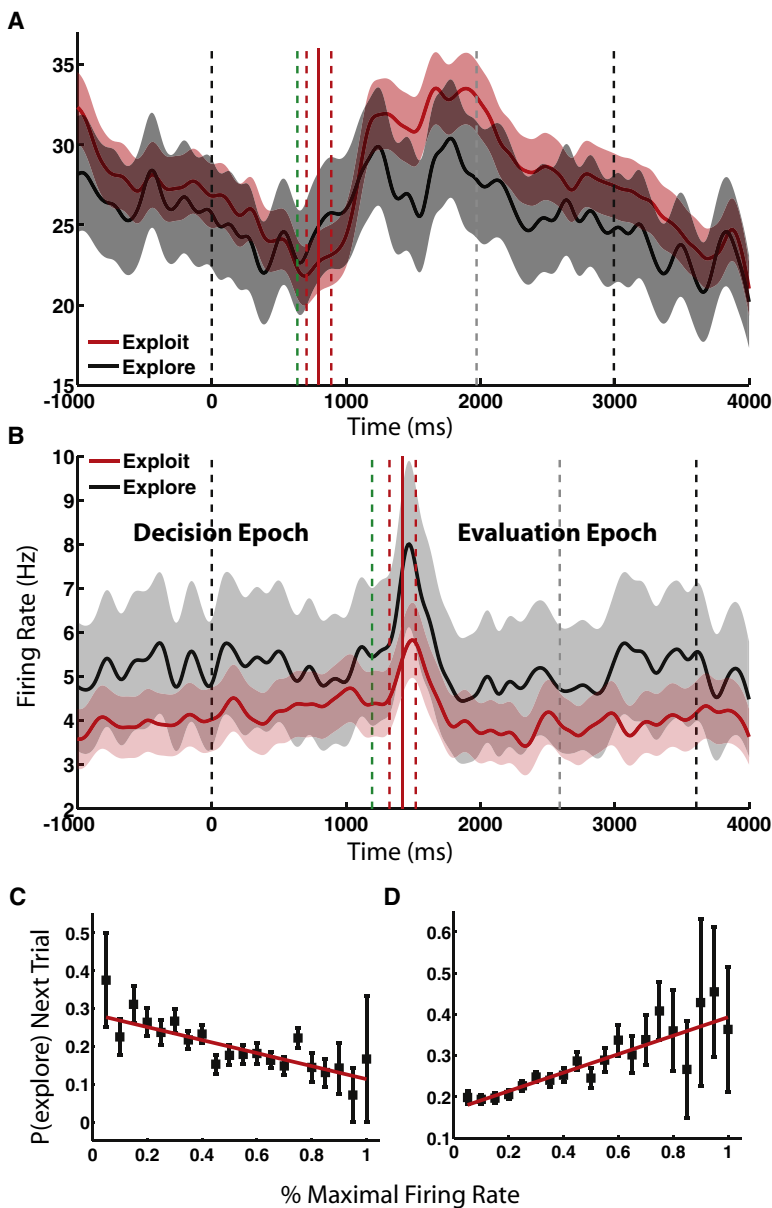


Figure 2. Neurons in CGp Distinguish between Exploration and Exploitation

(A and B) Sample neurons (neurons 62 and 12) with significant differences in firing during explore and exploit trials for both the decision and evaluation epochs, with firing aligned to trial onset. Individual cells might prefer either explore (black) or exploit (red) trials. The task begins at time 0. Onset of “go” cue (dashed green line), reward delivery (solid red line), beginning of intertrial interval (dashed gray line), and end of trial (second dashed black line) are mean times. Dashed red lines indicate \pm one standard deviation in reward onset. Shaded areas represent standard error of the mean (SEM). (C and D) Neurons in CGp encode probability of exploring on the next trial. Points are probabilities of exploring next trial as a function of percent maximal firing rate in the decision epoch, averaged over negatively and positively tuned populations of neurons (C and D, respectively). Error bars represent SEM.

a result, our task design allows for the possibility of disambiguating these phenomena.

Indeed, a partial correlation analysis of neuronal firing rates and upcoming decisions (firing rate in DE for decision in current trial; firing rate in EE for decision in following trial) that controlled for the effects of reward tuning, spatial tuning, target switching, and previous explore/exploit decision revealed significant correlations in 12% of neurons [$n = 11$, DE; $n = 10$, EE; $p < 0.05$ Spearman partial correlation; $p(n > 9) < 0.01$, binomial test]. Thus, even when all known effects on firing rate were accounted for, a significant number of neurons still exhibited clear predictive correlations with upcoming strategy. Collectively, these results indicate not only that single neurons in CGp receive information about both previous rewards and previous choices [19] and maintain that information across trials [11, 19], as reported previously, but that these same neurons also carry signals related to dynamic changes in choice strategy in a multiplexed format.

Discussion

We found that, when choosing among multiple targets whose relative values changed dynamically, neurons in posterior cingulate cortex signaled the distinction between trials on which monkeys pursued an exploratory rather than an exploitative strategy. This signal was robust against classifications of trials based on differing models of behavior, including a perfectly greedy strategy and a simple heuristic based on comparison to a reward threshold (see [Supplemental Data](#)). More importantly, single neurons signaled in graded fashion the probability of pursuing each strategy on upcoming trials.

Previous work has shown that CGp neurons are sensitive to reward [18], risk [19], and option switching [11] and integrate this information across multiple trials [11], but the present study generalizes the decision environment to one in which exploration and exploitation are distinguishable from a simple “win-stay/lose-shift” heuristic [11, 22] based only on the most recent reward received and in which outcomes must be evaluated in light of multiple options with dynamically changing rewards. As a result, relatively bad outcomes in rich environments might be acceptable under circumstances where all

sensitivity to strategy even when we controlled for the value of the reward the monkeys received.

Two other possible confounds arise from the known spatial tuning of CGp and the close relationship between exploration and simply switching between targets. As reported previously [21], we found that 63% of neurons were tuned for the location of the target chosen ($n = 52/83$, $p < 0.05$, one-way analysis of variance of mean firing rates for each target over all trials; see [Supplemental Data](#)). Across the population, 39% of neurons were significantly tuned for both reward size and target location [18], whereas 23% were tuned for neither (EE; 34% and 24%, respectively, in DE). However, the population as a whole showed no consistent target tuning across trials ($p > 0.2$, one-sample t test for contralateral and upper-hemifield tuning indices). In the case of target switching, because repeatedly choosing a poor target is not necessarily exploitative (if higher reward has recently been sampled elsewhere), there is not a strict one-to-one correspondence between exploitation and perseveration or between exploration and switching. As

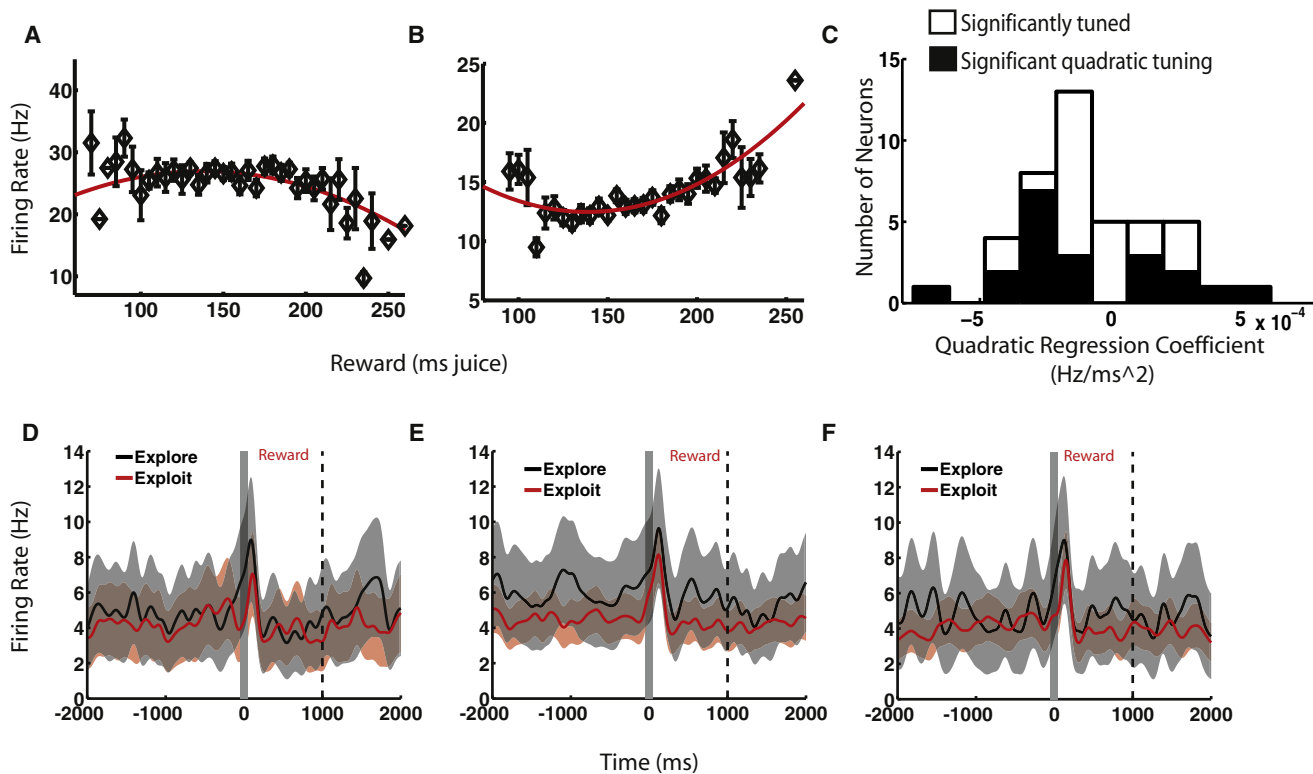


Figure 3. Reward Sensitivity of CGp Neurons Does Not Explain Sensitivity to Strategy

(A and B) Sample tuning curves from single neurons, showing firing as a function of reward received. Neurons display both positively and negatively monotonic tuning, as well as parabolic tuning (both concave up and concave down). (C) Histogram of quadratic regression coefficients. Positive values are concave up, as in (B). White bars represent neurons with significant overall regressions, as determined by F test. Black bars represent the subset of this population with significant coefficients for quadratic tuning. (D–F) Sample neuron (cell 12) firing rates on explore and exploit trials, controlled for low, medium, and high reward and aligned to reward onset. Bins are defined as follows: low, 115–135 ms; medium, 140–160 ms; high, 165–185 ms. In (D), the low case, the difference in mean reward sizes between explore and exploit conditions is 5 ms; in (E) and (F), the difference is less than 2 ms. Shaded regions represent SEM.

alternatives are poor and searching for better options necessitates choosing among several competing alternatives, each with a distinct reward history. Strategic decisions in such an environment thus require greater abstraction and integration of information than in comparatively static contexts because no single variable (or single trial) contains sufficient information on which to base a decision.

Together, these results invite the hypothesis that CGp is part of a network that monitors the outcomes of individual decisions and integrates that information into higher-level strategies spanning multiple choices. Thus, although the prolonged time courses of firing-rate changes in CGp are unlikely to be responsible for decisions on individual trials, their tonic activity levels might be responsible for encoding the gradual accumulation of information that gives rise to changes in strategy. However, we do not expect that this integration of single-trial history with strategic information is limited to CGp, nor is this its sole function. For example, Daw et al. [5] found that exploratory decisions are associated with activation in frontal polar cortex and intraparietal sulcus, whereas exploitative decisions are associated with activity in striatum. We speculate that information about individual rewards and reward predictions is initially computed in striatum, orbitofrontal cortex, and medial prefrontal cortex; subsequently combined with recent reward outcome and choice history and maintained online in CGp; and finally passed to the anterior

cingulate cortex (ACC), where it is utilized in the selection of appropriate actions. Moreover, reciprocal connections between ACC and CGp might play a role in learning which combinations of single-trial variables are most relevant when deciding among strategies to maximize reward.

In this framework, recent reward history, computational difficulty, stimulus novelty, memory load, and the statistics of the environment are distilled into a small number of task-related decision variables for the purposes of encoding and selecting among potential actions. Thus, individual neurons that compose this network would be expected to display sensitivity to the many single-trial variables like risk, reward, and spatial location that serve as its inputs. We found that these variables are represented multimodally by neurons in CGp. As a result, we suspect that CGp might play a key role in the process of learning the combinations of stimuli and accumulated statistics most relevant to making decisions, analogous to its role in simple conditioning [23–25]. This is in keeping with our observation of increased firing rates in response to block boundaries, near-threshold decisions, and aborted trials in the bandit task, as also observed in ACC [12]. If so, CGp dysfunction might be related to deficiencies in memory-guided learning and action selection observed in disorders like Alzheimer’s disease and obsessive-compulsive disorder, and its proper function might be crucial to the flexible adaptation of strategy in response to changing environments.

Experimental Procedures

Surgical Procedures

All procedures were approved by the Duke University Institutional Animal Care and Use Committee and were conducted in compliance with the Public Health Service's Guide for the Care and Use of Animals. Two rhesus monkeys (*Macaca mulatta*) served as test subjects for recording. A small prosthesis and a stainless steel recording chamber were attached to the calvarium. The chamber was placed over CGp at the intersection of the interaural and midsagittal planes. Animals were habituated to laboratory conditions and trained to perform oculomotor tasks for liquid reward. Animals received analgesics and antibiotics after all surgeries. The chamber was kept sterile with antibiotic washes and sealed with sterile caps.

Behavioral Techniques

Monkeys were familiar with the task. Eye position was sampled at 1000 Hz (camera, SR Research). Data were recorded by a computer running MATLAB (The Mathworks) with Psychtoolbox [26] and Eyelink [27]. Visual stimuli were squares (6° wide) on a computer monitor 50 cm away. A solenoid valve controlled juice delivery. Juice flavor was the same for each target.

On every trial, a central cue appeared and stayed on until the monkey fixated it. Fixation was maintained within a 1°–2° window. After a brief delay, the central cue disappeared and the four targets were displayed in the corners of the screen. Targets appeared in the same location each trial. After selection of a target, its border was illuminated and reward was delivered, followed by a 1 s intertrial interval. Rewards varied from 40 ms to 280 ms of solenoid open time in 5 ms increments (50–350 μ l, in 7.5 μ l increments). Juice volumes were linear in solenoid open time, and we have previously shown that monkeys discriminate juice volumes as small as 20 μ l [18]. All target values began at 200 μ l and reset each block. Blocks were 60 trials long and were cued by the appearance of a gray square in the center of the screen. Reward values for all targets changed each trial according to a biased random walk (see Supplemental Data).

Microelectrode Recording Techniques

Single electrodes (Frederick Haer Co.) were lowered under microdrive guidance (Kopf) until the waveforms of one to three individual neurons were isolated. Individual action potentials were identified by standard criteria and isolated on a Plexon system. Neurons were selected on the basis of the quality of isolation, but not on selectivity for the task. Recordings were made in areas 23 and 31 in the cingulate gyrus and ventral bank of the cingulate sulcus, anterior to the intersection of the marginal and horizontal rami.

Statistical Methods

We used an alpha of 0.05 as a criterion for significance. Peristimulus time histograms (PSTHs) were constructed by aligning spikes to trial events, averaging across trials, and smoothing by a Gaussian filter with 50 ms standard deviation. Shaded regions in PSTHs represent the standard error of the mean (\pm SEM), also Gaussian smoothed. Firing-rate modulation indices were calculated in each epoch as $m = (f_{\text{explore}} - f_{\text{exploit}}) / (f_{\text{explore}} + f_{\text{exploit}})$, where f is the firing rate averaged over the relevant subset of trials. Behavioral parameters were fitted by custom scripts written with the MATLAB Optimization Toolbox (The Mathworks). Details of modeling can be found in Supplemental Data.

Firing-Rate Analysis

Analyses utilized a binary, model-based classification of choices on each trial as exploratory or exploitative (see Supplemental Data). We tested for significant differences in firing rates in both the decision and evaluation epochs as a function of the explore/exploit classification of the decision on the current trial (that is, effects on DE_n and EE_n as a function of x_n , where x_n is the binary explore/exploit variable). We also tested for predictive correlations between firing in one epoch and upcoming decision (DE_n with x_n and EE_n with x_{n+1}). To do this, we binned firing rates for each neuron into deciles of percent maximal firing and examined the percentage of exploratory decisions made subsequent to epochs with firing rates in each bin. This allowed us to construct a probability of exploration as a function of percent maximal firing, which we averaged across significant cells of each tuning.

Our reward controls were performed by grouping the 45 distinct reward values into nine bins and comparing firing rates within each bin during the evaluation epoch on explore and exploit trials. Significance levels utilize a Bonferroni correction for the number of tests performed, which varied (not all bins contained an explore or exploit trial). Our reward-controlled

plots grouped the 15 middle rewards into three groups of five, denoted medium-high, medium-medium, and medium-low.

Our partial correlation analyses correlated (raw, unbinned) firing rate in a given epoch (DE_n or EE_n) with the upcoming explore/exploit decision (x_n and x_{n+1} , respectively). In each case, the correlation is controlled for spatial location (split into two variables, one for upper versus lower hemifield and one for left versus right hemifield, each taking values ± 1), previous received reward (r_{n-1} and r_n , respectively), chosen target switch (binary; s_n and s_{n+1} , respectively), and previous explore/exploit choice (x_{n-1} and x_n , respectively). Correlations were calculated as Spearman rank correlations and so allow for generic monotonic relations among variables.

Supplemental Data

Supplemental Data include Supplemental Experimental Procedures, six tables, and five figures and can be found with this article online at [http://www.cell.com/current-biology/supplemental/S0960-9822\(09\)01474-2](http://www.cell.com/current-biology/supplemental/S0960-9822(09)01474-2).

Acknowledgments

This work was supported by National Institute on Drug Abuse postdoctoral fellowship 023338-01 (B.Y.H.), National Institutes of Health grant R01EY013496 (M.L.P.), and the Duke Institute for Brain Studies (M.L.P.). We thank K. Watson for assistance in training the animals and A. Long for comments on the manuscript.

Received: April 3, 2009

Revised: July 8, 2009

Accepted: July 9, 2009

Published online: September 3, 2009

References

1. Stephens, D.W., and Krebs, J.R. (1986). *Foraging Theory* (Princeton, NJ: Princeton University Press).
2. Gittins, J.C. (1989). *Multi-armed Bandit Allocation Indices* (Chichester, NY: Wiley).
3. Whittle, P. (1988). Restless bandits: Activity allocation in a changing world. *J. Appl. Probab.* 25, 287–298.
4. Berry, D.A., and Fristedt, B. (1985). *Bandit Problems: Sequential Allocation of Experiments* (London: Chapman and Hall).
5. Daw, N.D., O'Doherty, J.P., Dayan, P., Seymour, B., and Dolan, R.J. (2006). Cortical substrates for exploratory decisions in humans. *Nature* 441, 876–879.
6. Wittmann, B.C., Daw, N.D., Seymour, B., and Dolan, R.J. (2008). Striatal activity underlies novelty-based choice in humans. *Neuron* 58, 967–973.
7. Walton, M.E., Devlin, J.T., and Rushworth, M.F. (2004). Interactions between decision making and performance monitoring within prefrontal cortex. *Nat. Neurosci.* 7, 1259–1265.
8. Yoshida, W., and Ishii, S. (2006). Resolution of uncertainty in prefrontal cortex. *Neuron* 50, 781–789.
9. Shima, K., and Tanji, J. (1998). Role for cingulate motor area cells in voluntary movement selection based on reward. *Science* 282, 1335–1338.
10. Kennerley, S.W., Walton, M.E., Behrens, T.E., Buckley, M.J., and Rushworth, M.F. (2006). Optimal decision making and the anterior cingulate cortex. *Nat. Neurosci.* 9, 940–947.
11. Hayden, B.Y., Nair, A.C., McCoy, A.N., and Platt, M.L. (2008). Posterior cingulate cortex mediates outcome-contingent allocation of behavior. *Neuron* 60, 19–25.
12. Quilodran, R., Rothe, M., and Procyk, E. (2008). Behavioral shifts and action valuation in the anterior cingulate cortex. *Neuron* 57, 314–325.
13. Rudebeck, P.H., Behrens, T.E., Kennerley, S.W., Baxter, M.G., Buckley, M.J., Walton, M.E., and Rushworth, M.F. (2008). Frontal cortex subregions play distinct roles in choices between actions and stimuli. *J. Neurosci.* 28, 13775–13785.
14. Kobayashi, Y., and Amaral, D.G. (2003). Macaque monkey retrosplenial cortex: II. Cortical afferents. *J. Comp. Neurol.* 466, 48–79.
15. Vogt, B.A., Finch, D.M., and Olson, C.R. (1992). Functional heterogeneity in cingulate cortex: The anterior executive and posterior evaluative regions. *Cereb. Cortex* 2, 435–443.
16. Vogt, B.A., and Gabriel, M. (1993). *Neurobiology of Cingulate Cortex and Limbic Thalamus: A Comprehensive Handbook* (Boston: Birkhauser).

17. Vogt, B.A., Gabriel, M., Vogt, L.J., Poremba, A., Jensen, E.L., Kubota, Y., and Kang, E. (1991). Muscarinic receptor binding increases in anterior thalamus and cingulate cortex during discriminative avoidance learning. *J. Neurosci.* *11*, 1508–1514.
18. McCoy, A.N., Crowley, J.C., Haghghian, G., Dean, H.L., and Platt, M.L. (2003). Saccade reward signals in posterior cingulate cortex. *Neuron* *40*, 1031–1040.
19. McCoy, A.N., and Platt, M.L. (2005). Risk-sensitive neurons in macaque posterior cingulate cortex. *Nat. Neurosci.* *8*, 1220–1227.
20. Dean, H.L., and Platt, M.L. (2006). Allocentric spatial referencing of neuronal activity in macaque posterior cingulate cortex. *J. Neurosci.* *26*, 1117–1127.
21. Dean, H.L., Crowley, J.C., and Platt, M.L. (2004). Visual and saccade-related activity in macaque posterior cingulate cortex. *J. Neurophysiol.* *92*, 3056–3068.
22. Barraclough, D.J., Conroy, M.L., and Lee, D. (2004). Prefrontal cortex and decision making in a mixed-strategy game. *Nat. Neurosci.* *7*, 404–410.
23. Gabriel, M., Foster, K., and Orona, E. (1980). Interaction of laminae of the cingulate cortex with the anteroventral thalamus during behavioral learning. *Science* *208*, 1050–1052.
24. Gabriel, M., Sparenborg, S.P., and Stolar, N. (1987). Hippocampal control of cingulate cortical and anterior thalamic information processing during learning in rabbits. *Exp. Brain Res.* *67*, 131–152.
25. Gabriel, M., Kubota, Y., Sparenborg, S., Straube, K., and Vogt, B.A. (1991). Effects of cingulate cortical lesions on avoidance learning and training-induced unit activity in rabbits. *Exp. Brain Res.* *86*, 585–600.
26. Brainard, D.H. (1997). The Psychophysics Toolbox. *Spat. Vis.* *10*, 433–436.
27. Cornelissen, F.W., Peters, E.M., and Palmer, J. (2002). The EyeLink Toolbox: Eye tracking with MATLAB and the Psychophysics Toolbox. *Behav. Res. Methods Instrum. Comput.* *34*, 613–617.