

Posterior Cingulate Cortex Mediates Outcome-Contingent Allocation of Behavior

Benjamin Y. Hayden,^{1,*} Amrita C. Nair,¹ Allison N. McCoy,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

*Correspondence: hayden@neuro.duke.edu

DOI 10.1016/j.neuron.2008.09.012

SUMMARY

Adaptive decision making requires selecting an action and then monitoring its consequences to improve future decisions. The neuronal mechanisms supporting action evaluation and subsequent behavioral modification, however, remain poorly understood. To investigate the contribution of posterior cingulate cortex (CGp) to these processes, we recorded activity of single neurons in monkeys performing a gambling task in which the reward outcome of each choice strongly influenced subsequent choices. We found that CGp neurons signaled reward outcomes in a nonlinear fashion and that outcome-contingent modulations in firing rate persisted into subsequent trials. Moreover, firing rate on any one trial predicted switching to the alternative option on the next trial. Finally, microstimulation in CGp following risky choices promoted a preference reversal for the safe option on the following trial. Collectively, these results demonstrate that CGp directly contributes to the evaluative processes that support dynamic changes in decision making in volatile environments.

INTRODUCTION

The brain mechanisms that monitor behavioral outcomes and subsequently update representations of action value remain obscure (Platt, 2002). Multiple brain areas have been implicated in outcome monitoring, particularly the anterior cingulate cortex (ACC; Walton et al., 2004; Matsumoto et al., 2007; Quilodran et al., 2008; Behrens et al., 2007; Kennerley et al., 2006; Rushworth et al., 2007; Shima and Tanji, 1998), while other areas, including lateral prefrontal cortex, parietal cortex, and dorsal striatum, have been linked to coding action value (Barraclough et al., 2004; Lau and Glimcher, 2007; Kobayashi et al., 2002; Leon and Shadlen, 1999; Sugrue et al., 2004; Watanabe, 1996). Posterior cingulate cortex (CGp) may serve as a link between these processes. CGp is reciprocally connected with both the ACC and parietal cortex (Kobayashi and Amaral, 2003). CGp neurons respond to rewards (McCoy et al., 2003), signal preferences in a gambling task with matched reward rates (McCoy and Platt,

2005), and signal omission of predicted rewards (McCoy et al., 2003). Collectively, these observations suggest the hypothesis that CGp contributes to the integration of actions and their outcomes and thereby influences subsequent changes in behavior.

To test this hypothesis, we studied the responses of single neurons, as well as the effects of microstimulation, in CGp in monkeys performing a gambling task. Monkeys prefer the risky option in this task, but their local pattern of choices strongly depends on the most recent reward obtained (Hayden and Platt, 2007; McCoy and Platt, 2005). This task is thus an ideal tool for studying the neural mechanisms underlying outcome monitoring and subsequent changes in choice behavior. We specifically probed how CGp neurons respond to gamble outcomes, how such signals influence future choice behavior, and whether artificial activation of CGp systematically perturbs impending decisions.

Consistent with our hypothesis, we found that reward outcomes influenced both neuronal activity and the future selection of action. Specifically, monkeys were more likely to switch to an alternative when they received less than the maximum reward obtainable, and CGp neurons responded most strongly to rewards that deviated from this maximum value. Moreover, reward outcome signals persisted into future trials and predicted subsequent changes in choice behavior. Finally, microstimulation following the resolution of risky gambles increased the probability that monkeys would switch to the safe option on the next trial. These results build on our prior findings that CGp neurons carry information correlated with preferences in a gamble and directly implicate this brain area in the neural processes that link reward outcomes to dynamic changes in behavior.

RESULTS

Reward Outcomes Influence Local Patterns of Choice

On each trial of the gambling task (Figure 1A), monkeys indicated their choice by shifting gaze to either a safe or a risky target. The safe target provided a reward of predictable size (200 μ l juice); the risky target yielded either a larger or smaller reward (varied randomly). The expected value of the risky option was equal to the safe option; reward variance was altered in blocks (see *Experimental Procedures* for details; McCoy and Platt, 2005).

Monkeys were risk seeking overall, but their local pattern of choices strongly depended on the previous reward. After receiving a large reward for selecting the risky option, monkeys were more likely to select it again than after receiving the smaller

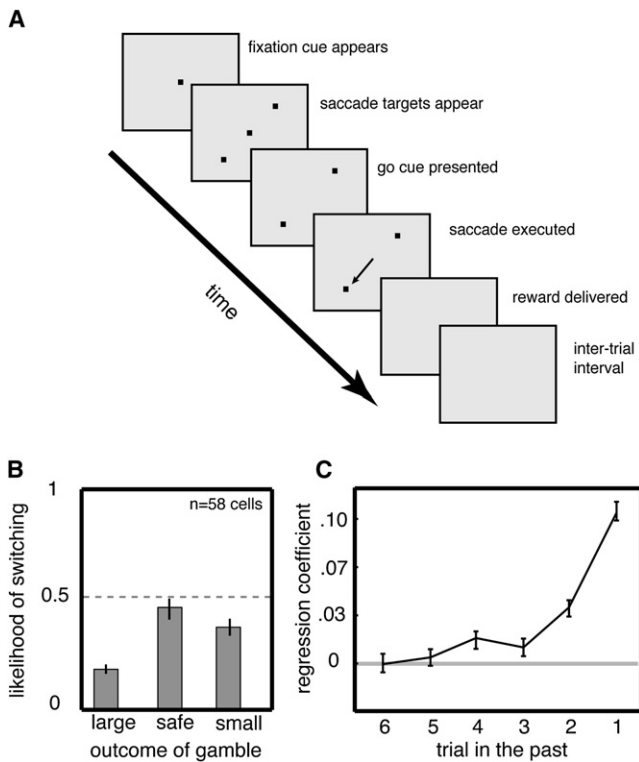


Figure 1. Monkeys' Choices Are a Nonlinear Function of Previous Reward Outcomes

(A) Gambling task. After monkey fixated on a central square, two eccentric response targets appeared. After a delay (1 s), the central square was extinguished, cuing the monkey to shift gaze to either of the two targets to receive a reward. One target (safe) offered a certain juice volume the other target offered an unpredictable (risky) volume.

(B) Average frequency of switching from risky to safe or vice versa in all recording sessions. Monkeys were more likely to choose the risky option again following a large reward than following a small reward. (C) Influence of the outcome of a single reward declines over time. Logistic regression coefficient for likelihood of choosing risky option as a function of reward outcome up to six trials in the past. Bars indicate one standard error.

reward (as reported previously, McCoy and Platt, 2005; Figure 1B). Monkeys were 83.3% likely to choose the risky option following a large reward and only 64.2% likely to choose the risky option following a small reward (Student's *t* test, $p < 0.001$, $n = 58$ sessions). The likelihood of switching was greater following a small than following a large reward in every individual session ($p < 0.05$, binomial test on outcomes of individual trials in each session). After choosing the safe option, the likelihood of switching to the risky option was 45.4%.

We next asked whether obtaining a large or small reward on any single trial influenced choices on trials beyond the next one (Figure 1C). We performed a logistic regression of choice (risky or safe) on the outcomes of the most recent six choices, an analysis that produces something akin to a behavioral kernel (cf. Lau and Glimcher, 2007; Sugrue et al., 2004). We found that, on average, each reward outcome influenced choices up to about four trials in the future, but the influence of any single reward diminished rapidly across trials. These results demonstrate

that monkeys adjust their choices in this task by integrating recent reward outcomes and comparing this value with the largest—and most desirable—reward. When these two values are similar, monkeys tend to repeat the last choice, but when they differ, monkeys explore the alternative option.

CGp Neurons Signal Reward Outcomes and Maintain This Information across Trials

In a prior study, we showed that neuronal activity in CGp tracks both risk level and behavioral preference for the risky option; the current report complements these previous findings by probing how neuronal activity in CGp contributes to the neural processes that link specific reward outcomes to individual decisions. To do this, we analyzed the activity of 58 neurons in 2 monkeys (32 in monkey N and 26 in monkey B; data from 42 of these was analyzed in McCoy and Platt [2005]). Figure 2A shows the response of a single neuron on trials in which the monkey chose the risky option and received a large reward (dark gray line) or the small reward (black line). Responses are aligned to reward offset (time = 0). After the reward, neuronal activity on small reward (and safe) trials was greater than on large reward trials (Student's *t* test, $p < 0.01$). In a 1 s epoch following reward, most neurons showed significant differences for large and small rewards (Figure 2C, 74%, $n = 43/58$, $p < 0.05$, Student's *t* test on individual trials within each neuron). Most neurons increased firing following small rewards ($n = 28$, 48% of all neurons) while a minority ($n = 15$, 26%) increased firing following large rewards. Across all neurons, firing rates were greater after small rewards than after large rewards (0.98 sp/s in all neurons, $p = 0.01$, Student's *t* test; 1.53 sp/s in significantly modulated neurons; Figure 2B). Thus, CGp neurons preferentially responded to reward outcomes that deviated from the maximum obtainable reward.

To directly influence subsequent changes in behavior, neuronal responses to rewards must persist across delays between trials. We therefore examined the effect of reward outcome on neuronal activity at the beginning of the next trial (500 ms epoch preceding the fixation cue, time = 0 in Figure 2D). Figure 2D shows the average firing rate of a sample neuron (same as Figure 2A) on trials following large (dark gray line) and small (black line) rewards. Neuronal responses were significantly greater following small rewards than following large rewards (Student's *t* test, $p < 0.01$, $n = 293$ trials). Many neurons (34%, $n = 20/58$) showed a significant change in firing rate reflecting reward outcome on the previous trial ($p < 0.05$, Student's *t* test on individual trials). The average difference in firing rate at the beginning of the next trial (0.45 sp/s in all cells, 0.85 sp/s in significantly modulated cells) was smaller than the average difference in firing rate immediately following the reward (student's *t* test, $p < 0.001$). Across all neurons, the firing rate change after the outcome of a gamble was correlated with the firing rate of the same neuron at the beginning of the next trial (Figure 2E, $r^2 = 0.4811$, $p < 0.001$, correlation test, $n = 58$ neurons). These modulations were also present during the 0.5 s epoch preceding the saccade—41% of neurons ($n = 24/58$) showed a significant change in firing rate that reflected the outcome of the previous trial. The average size of the response change during this epoch was similar to that observed preceding the trial (0.52 sp/s). Collectively, these results indicate that CGp neurons maintain

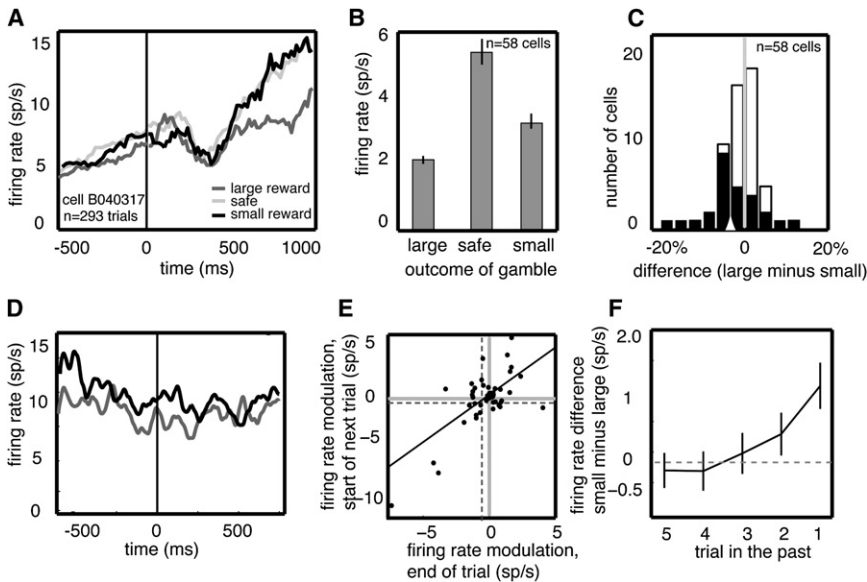


Figure 2. CGp Neurons Signal Reward Outcomes in a Nonlinear Fashion and Maintain This Information for Several Trials

(A) PSTH for example neuron following reward delivery, aligned to reward offset. Responses were significantly greater following small or medium reward than following large reward.

(B) Bar graph showing the average response of all neurons in the population to large, medium, and small rewards. Responses are averaged over a 1 s epoch beginning at the time of reward offset ($t = 0$). Bars indicate one standard error.

(C) Histogram of reward modulation indices. Index was defined as the response difference to large and small rewards. Negative values correspond to neurons for which small rewards evoked greater responses than larger rewards. Most neurons responded more following small rewards than following large rewards. Black and white bars represent significantly and nonsignificantly modulated cells, respectively.

(D) Average responses of example neuron at the beginning of a trial (the 500 ms before the fixation cue that began the trial, $t = 0$) following a trial in

which a small (black) or large (dark gray) reward was received. Responses are aligned to the beginning of the trial.

(E) Scatter plot showing the average firing rate modulation as a function of trial outcome during a 1 s epoch following reward on a given trial (x axis) and a 0.5 s epoch following acquisition of fixation on the subsequent trial (y axis) for each neuron in the population. Firing rates were correlated in these two epochs, and the average size of the outcome modulation immediately following the gamble (vertical dashed line) was greater than the average size of the outcome modulation on the subsequent trial (horizontal dashed line).

(F) Average effect of reward outcome on neuronal activity up to five trials in the future. Bars indicate one standard error.

information about reward outcomes from one trial to the next. Because representation of such information is critical for many forms of reward-based decision making, irrespective of whether the options are presented in the form of a gamble, the present results demonstrate the importance of CGp for action-outcome learning in general.

To determine whether the influence of reward outcomes persisted for multiple trials, we performed a multiple linear regression of neuronal firing rates on the outcomes of the most recent five gambles (Figure 2F). Neuronal firing rates during the 1 s epoch following reward offset reflected the outcome of the most recent two gambles, with weaker influence of the second-to-last trial than the most recent. The influence of a single reward outcome was approximately 1.2 sp/s (11.3% of average firing rate) on the next trial and 0.42 sp/s (3.9% of average firing rate) two trials in the future. Thus, CGp neurons maintain reward information across multiple trials—a delay of several seconds—and this influence diminishes with time at a rate similar to the diminishing influence of reward outcomes on behavior.

CGp Neurons Signal Future Changes in Choice Behavior

Persistent reward-related activity in CGp may signal the need to switch from the previously selected option to an alternative. To test this hypothesis, we compared neuronal responses on trials preceding a switch from risky to safe or vice versa with trials that did not precede such a switch. To control for the correlation between reward outcomes and changes in neuronal activity, we performed this analysis for each possible outcome sepa-

rately (large, small, safe). Figure 3A shows PSTHs for an example neuron (different from the one shown above) on trials preceding a switch (gray line) and on trials that did not precede a switch (black line). Neuronal activity predicted the subsequent switch, even before the resolution of the gamble.

Responses of many neurons (36%, $n = 21/58$) during the 1 s epoch following reward offset were modulated prior to subsequent preference reversals (Figure 3C). For the majority of these (24%, $n = 14$, 67% of significantly modulated neurons), firing increased before a change, and in the minority (12%, $n = 7$, 33% of significantly modulated neurons), firing decreased. Average firing rate of all neurons was greater before a switch (Figure 3B, 0.54 sp/s, 2.37 sp/s in significantly modulated cells, $p = 0.03$, Student's *t* test, $n = 58$ cells). A 2×3 ANOVA on population responses confirmed that there were significant main effects of impending switches (switch and no switch, $p < 0.04$) and trial outcome (large, medium, small, $p < 0.01$) on firing rates in the postreward epoch, but no significant interactions ($p > 0.5$). The same ANOVA applied to the activity of individual neurons revealed a significant statistical interaction between impending switch and trial outcome in only 8.6% of neurons ($n = 5/58$, $p < 0.05$ for result of ANOVA on individual trials).

We next compared the size of the switch-related modulation in early, middle, and late epochs during each trial (corresponding to -2 to -1 , -1 to 0 , and 0 to 1 on Figure 3A). We used a 2×3 ANOVA for firing rates against switch (switch or no switch) and epoch (early, middle, and late). We found no main effect of epoch ($p = 0.4$). We note that the lack of an effect here may

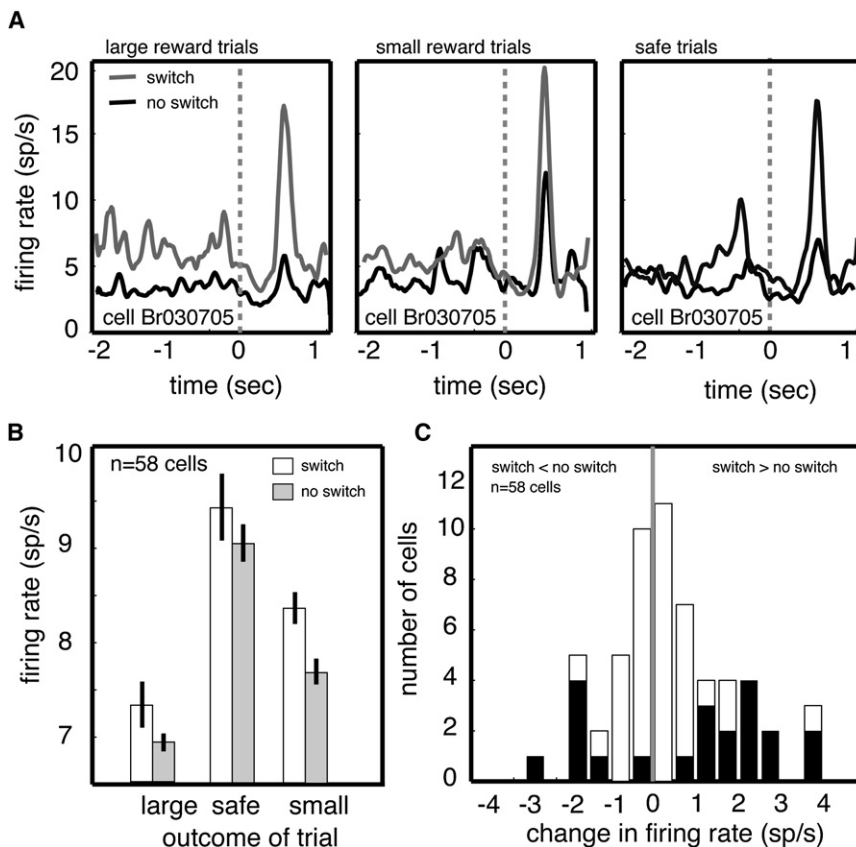


Figure 3. CGp Neurons Signal Subsequent Changes in Choice Behavior

(A) Responses of a single neuron on trials immediately preceding a switch in choice behavior (gray line) and no switch in behavior (black line) for different reward outcomes. PSTHs are aligned to the time at which the gamble is resolved (indicated as time zero on the graph).

(B) Bar graph showing the average firing rate during a 1 s epoch beginning at reward offset for all neurons in the population on trials that preceded switches (white bars) and that did not (gray bars). Responses were greater preceding a switch, although this effect was only significant following choices of the risky option. Bars indicate one standard error.

(C) Histogram of firing rate modulations associated with switching behavior (neurons with significant modulations are indicated by black bars).

reflect insufficient data for analysis. These data suggest that switching related changes in activity are not restricted to the end of the trial and hint that local choice patterns may reflect the integrated outcome of ongoing evaluative processes. In a previous study, we showed that firing rates of CGp neurons predict choices of the risky option in a gamble (McCoy and Platt, 2005); the present results show that this decision is influenced by the outcomes of recent trials and that this information is represented in a persistent fashion by neuronal activity in CGp.

Microstimulation in CGp Promotes Behavioral Switching

To demonstrate a causal role for CGp in choice behavior, we examined the effects of postreward microstimulation. Because this epoch showed reward-dependent modulation in firing that predicted impending changes in choice behavior, we hypothesized that stimulation at this time would increase the likelihood of switching to the alternative option. We examined the choices of two monkeys performing the risky decision-making task in separate sessions. Monkey B provided data for the experiments described above; Monkey S was familiar with the task but did not contribute other data. Timing (reward offset) and duration (500 ms) of stimulation were chosen to approximate the timing and duration of the neuronal response to rewards. Current was 200 μ A.

Stimulation following rewards delivered for risky choices increased the frequency of switching by 3.14% (Student's t test, $p = 0.021$; Figure 4A). Stimulation did not affect behavior following

choices of the safe option (effect size was -0.075% , $p = 0.57$, Student's t test). The effect of stimulation on switching following large rewards (5.2%) was significantly ($p < 0.05$, Student's t test) greater than the effect of stimulation on switching following small rewards (2.6%). Thus, stimulation primarily induced monkeys to switch to the safe option after receiving a large reward—precisely those trials on which firing was lowest and on which monkeys were most loathe to switch, in the absence of stimulation. An ANOVA confirmed a main effect of risky versus safe choice ($p = 0.037$), but no effect of side (ipsi- or contralateral, $p > 0.05$), and no interaction ($p > 0.05$). The failure to observe behavioral changes following safe choices may be due to monkey's general dislike for, and thus high baseline rate of switching away from, this option. Importantly, the differential effects of stimulation following risky and safe choices preclude the possibility that the observed behavioral effects were simply due to stimulation-evoked discomfort, distraction, avoidance, or random guessing.

One caveat is that stimulation could have introduced a motor bias or directly evoked saccades. The ~ 3 s delay between stimulation (at the end of one trial) and the subsequent choice militates against this possibility, as does the lack of a side bias in the stimulation effect. Nonetheless, we performed a control experiment to determine whether stimulation evokes saccades. In 20 sessions, we performed un signaled stimulation (same parameters as above) for 200 trials in the absence of a task (Figures 4B and 4C). We found no effect of stimulation on either eye position or velocity for 2 s following stimulation (Student's t test, $p > 0.5$ for each individual session and $p > 0.5$ for all sessions together). Mean eye position did not differ during a 500 ms epoch before stimulation and during any of four subsequent 500 ms epochs (Student's t test, $p > 0.5$ in all cases). Likewise, eye velocity did not vary between the 500 ms epoch preceding stimulation and any of four subsequent 500 ms epochs.

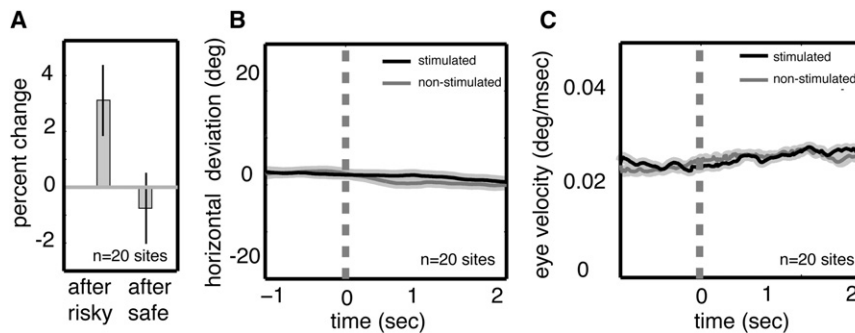


Figure 4. Microstimulation in CGp Increases the Frequency of Local Preference Reversals

(A) Vertical bars indicate change in likelihood of switching to the other option when stimulation occurs following a reward. Stimulation leads to a significantly increased likelihood of switching to the safe option following risky choices, but not after safe choices. Black bars indicate one standard error of the mean.

(B and C) Results of control experiment showing the stimulation does not evoke changes in eye velocity or position. Traces are aligned to time of stimulation in the absence of a task. Shaded regions indicate one standard error.

DISCUSSION

We found that firing rates of CGp neurons nonlinearly signaled reward outcomes in a gambling task. Most neurons fired more for small rewards than for large rewards and fired most strongly for medium-sized but predictable rewards. Neuronal responses thus paralleled the effects of reward outcomes on subsequent choices. This nonlinear relationship between neuronal responses during dynamic decision making contrasts with the monotonic relationship observed during imperative orienting (McCoy et al., 2003). Moreover, reward outcome signals persisted across the delays between trials and predicted impending changes in choice behavior. We confirmed a causal role for CGp in outcome evaluation and behavioral adjustment by showing that microstimulation following the most desired reward outcome increased the likelihood of exploring the alternative on the next trial. We acknowledge that such an account is likely to be overly simplistic and is thus only a first step. For example, recent research indicates that neurons in dLPFC maintain information about the outcomes of trials and reflect changes in behavior as well (Seo and Lee, 2007) and thus must be incorporated into any model of decision making.

Notably, information about reward outcomes was maintained by relatively slow, long-lasting changes in firing rate in CGp. Such changes are consistent with the generally tonic changes in firing rate of these cells in a variety of tasks (McCoy et al., 2003; McCoy and Platt, 2005), and thus we conjecture that they may reflect a working memory or attention-related process. Indeed, data from neuroimaging (Maddock et al., 2003) and lesion (Gabriel, 1990) studies strongly implicate CGp in working memory, attention, and general arousal (Raichle and Mintun, 2006)—all processes that contribute to adaptive decision making.

The current results complement our previous study of CGp neurons using the same gambling task in several important ways (McCoy and Platt, 2005). In the previous study, we focused on neuronal activity occurring prior to and after the expression of choice, whereas the present work focuses on neuronal activity following reward and immediately preceding ensuing trials. In our previous study, we demonstrated that neuronal activity in CGp is correlated with both the amount of risk associated with an option and monkeys' proclivity to choose it. That work thus linked CGp to the subjective aspects of decision making. By contrast, the present study demonstrates that CGp neurons

signal decision outcomes in a nonlinear fashion, maintain this information in a buffer between trials, and predict future changes in behavior. Moreover, microstimulation in CGp promotes exploration of the previously antipreferred option. Together, these new findings show that CGp directly contributes to the neural processes that evaluate reward outcomes in subjective terms and use this information to influence subsequent decisions. Although these observations were made in the context of risky decisions, they apply equally well to any action that must be evaluated in order to make better decisions in the future. Taken together, the findings of these two studies highlight the dynamic nature of information processing in CGp, which is not restricted to any single epoch or aspect of task performance but instead continuously adapts to changes in both the external environment and internal milieu.

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 subjects for recording; another served as a subject for microstimulation. A small prosthesis and a stainless steel recording chamber were attached to the calvarium, and a filament for ocular monitoring was implanted using standard techniques. The chamber was placed over CGp at the intersection of the interaural and midsagittal planes. Animals were habituated to training 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 500 Hz (scleral coil, Riverbend Instruments) or 1000 Hz (camera, SR Research). Data were recorded by a computer running Gramalkyn (www.ryklinsoftware.com) or Matlab (Mathworks) with Psychtoolbox (Brainard, 1997) and EyeLink Toolbox (Cornelissen et al., 2002). Visual stimuli were LEDs (LEDtronics) on a tangent screen 145 cm (57 inches) from the animal, or squares (2° wide) on a computer monitor 50 cm away. A solenoid valve controlled juice delivery.

On every trial, a central cue appeared and stayed on until the monkey fixated it. Fixation was maintained within a 1° window (in a small fraction of sessions, we used a 2° window). After a brief delay, two eccentric targets appeared while the cue remained illuminated (the decision period); then the central target disappeared and the animal shifted gaze to one of the two eccentric targets. The targets were placed so that one was within the neuron's response field while the other was located 180° away. Failure to saccade led to the immediate

end of the trial and a 5 s timeout period. The safe target offered 200 μ l juice; the risky target offered one of two rewards, selected at random for each block and not signaled. The average reward for the risky target was the value of the safe target. The size of the risky rewards from drawn from the following list (μ l): 33/367, 67/333, 133/267, 167/233, 180/220, 187/213.

The locations of the safe and risky targets and the variance of the risky target were varied in blocks of 50 trials for the first set of 42 neurons, and in blocks of 10 trials for the remainder. We observed no systematic differences in these two sets of neurons, so we combined the data for analysis. In some sessions, a 300 ms white noise signal provided a secondary reinforcer; this did not significantly affect behavior. To ensure that reward volume was a linear function of solenoid open time, we performed calibrations before, during, and after both experiments. We found that the relationship between open time and volume was linear and did not vary from day to day or month to month. Following reward, the monitor was left blank for 2 s in some sessions and 4 s in others.

Microelectrode Recording Techniques

Single electrodes (Frederick Haer Co) were lowered by microdrive (Kopf) until the waveform of a single neuron was isolated. Individual action potentials were identified in hardware by time and amplitude criteria (BAK electronics). Neurons were selected on the basis of the quality of isolation and sometimes by saccadic responses, but not on selectivity for the gambling task. We confirmed the location of the electrode using a hand-held digital ultrasound device (Sonosite 180) placed against the recording chamber; 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. A subset of the neurons analyzed in this study were also analyzed in a previous study ($n = 42$); others were collected specifically for this study ($n = 16$).

Stimulation

Stimulation was performed in a separate set of sessions using the same task as above. One monkey was the same as used in the previous study (B); the other was not used in the previous study (S). S's behavior matched that of the other two monkeys. For stimulation, CGp was identified by stereotaxis, at depths where most neurons had been obtained. Pulses were generated using a Master-8 Pulse Generator (A.M.P.I.) and converted to constant current using a BP isolator (Frederick Haer Co.). Stimulation began at the time of the reward and lasted 500 ms. This time window overlaps with the period of strongest evoked activity. Current was 200 μ A, delivered at 200 Hz; each biphasic pulse lasted 200 μ s. In a control experiment, stimulation occurred at random time, no more than once every 30 s, during a long break period. No external signal predicted stimulation. We collected 200 trials in each session, 10 sessions with monkey S, 10 with monkey B.

Statistical Methods

We used an alpha of 0.05 as a criterion for significance. Dependence of choices on task factors was estimated using linear or logistic regression. Peristimulus time histograms (PSTHs) were constructed by aligning spikes to trial events, averaging across trials, and smoothing with a 100 ms boxcar. In most analyses, data were aligned to the time at which reward delivery ended. To examine the firing rate at the beginning of the trial, we aligned to the onset of the first fixational cue. Statistics were performed on binned firing rates (1 s or 0.5 s). To compare firing rates across trials for single neurons, tests were performed on individual trials; to compare firing rates across neurons, tests were performed on average rates for individual neurons. Standard errors were standard error of mean firing rates across the entire bin. In all cases, the results of *t* tests were confirmed with a bootstrap (i.e., nonparametric) hypothesis test. We compared the difference between the means of the two distributions to those of 10,000 randomized distributions. A hypothesis was accepted as significant if the observed difference occurred within the lowest 250 or the highest 250 randomized differences, providing a two-tailed alpha of 0.05.

SUPPLEMENTAL DATA

The Supplemental Data include one figure, Supplemental Results, and Supplemental Experimental Procedures and can be found with this article online at [http://www.neuron.org/supplemental/S0896-6273\(08\)00760-5](http://www.neuron.org/supplemental/S0896-6273(08)00760-5).

ACKNOWLEDGMENTS

We thank Vinod Venkatraman for help in developing Psychtoolbox for our rig. We thank Sarah Heilbronner for comments on the manuscript. This work was supported by the NIH EY013496 (M.L.P.) and a Kirschstein NRSA DA023338 (B.Y.H.).

Accepted: September 4, 2008

Published: October 8, 2008

REFERENCES

- 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.
- Behrens, T.E., Woolrich, M.W., Walton, M.E., and Rushworth, M.F. (2007). Learning the value of information in an uncertain world. *Nat. Neurosci.* 10, 1214–1221.
- Brainard, D.H. (1997). The psychophysics toolbox. *Spat. Vis.* 10, 433–436.
- Cornelissen, F.W., Peters, E., and Palmer, J. (2002). The eyelink toolbox: eye tracking with MATLAB and the psychophysics toolbox. *Behav. Res. Methods Instrum. Comput.* 34, 613–617.
- Gabriel, M. (1990). Functions of anterior and posterior cingulate cortex during avoidance learning in rabbits. *Prog. Brain Res.* 85, 467–482.
- Hayden, B.Y., and Platt, M.L. (2007). Temporal discounting predicts risk sensitivity in rhesus macaques. *Curr. Biol.* 17, 49–53.
- 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.
- Kobayashi, Y., and Amaral, D.G. (2003). Macaque monkey retrosplenial cortex: II. Cortical afferents. *J. Comp. Neurol.* 466, 48–79.
- Kobayashi, S., Lauwereyns, J., Koizumi, M., Sakagami, M., and Hikosaka, O. (2002). Influence of reward expectation on visuospatial processing in macaque lateral prefrontal cortex. *J. Neurophysiol.* 87, 1488–1498.
- Lau, B., and Glimcher, P.W. (2007). Action and outcome encoding in the primate caudate nucleus. *J. Neurosci.* 27, 14502–14514.
- Leon, M.I., and Shadlen, M.N. (1999). Effect of expected reward magnitude on the response of neurons in the dorsolateral prefrontal cortex of the macaque. *Neuron* 24, 415–425.
- Maddock, R.J., Garrett, A.S., and Buonocore, M.H. (2003). Posterior cingulate activation by emotional words: fMRI evidence from a valence decision task. *Hum. Brain Mapp.* 18, 30–41.
- Matsumoto, M., Matsumoto, K., Abe, H., and Tanaka, K. (2007). Medial prefrontal cell activity signaling prediction errors of action values. *Nat. Neurosci.* 10, 647–656.
- McCoy, A.N., and Platt, M.L. (2005). Risk-sensitive neurons in macaque posterior cingulate cortex. *Nat. Neurosci.* 8, 1220–1227.
- McCoy, A.N., Crowley, J.C., Haghigian, G., Dean, H.L., and Platt, M.L. (2003). Saccade reward signals in posterior cingulate cortex. *Neuron* 40, 1031–1040.
- Platt, M.L. (2002). Neural correlates of decisions. *Curr. Opin. Neurobiol.* 12, 141–148.
- Quilodran, R., Rothe, M., and Procyk, E. (2008). Behavioral shifts and action valuation in the anterior cingulate cortex. *Neuron* 57, 314–325.
- Raichle, M.E., and Mintun, M.A. (2006). Brain work and brain imaging. *Annu. Rev. Neurosci.* 29, 449–476.

Rushworth, M.F., Behrens, T.E., Rudebeck, P.H., and Walton, M.E. (2007). Contrasting roles for cingulate and orbitofrontal cortex in decisions and social behaviour. *Trends Cogn. Sci.* 11, 168–176.

Seo, H., and 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., and Tanji, J. (1998). Role for cingulate motor area cells in voluntary movement selection based on reward. *Science* 282, 1335–1338.

Sugrue, L.P., Corrado, G.S., and Newsome, W.T. (2004). Matching behavior and the representation of value in the parietal cortex. *Science* 304, 1782–1787.

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.

Watanabe, M. (1996). Reward expectancy in primate prefrontal neurons. *Nature* 382, 629–632.