



Published in final edited form as:

Ann N Y Acad Sci. 2007 May ; 1104: 21–34.

Should I Stay or Should I Go?:

Transformation of Time-Discounted Rewards in Orbitofrontal Cortex and Associated Brain Circuits

MATTHEW R. ROESCH^a, DONNA J. CALU^b, KATHRYN A. BURKE^b, and GEOFFREY SCHOENBAUM^{a,b,c}

^a*Department of Anatomy and Neurobiology, University of Maryland School of Medicine, Baltimore, Maryland, USA*

^b*Program in Neuroscience, University of Maryland School of Medicine, Baltimore, Maryland, USA*

^c*Department of Psychology, University of Maryland Baltimore County, Baltimore, Maryland, USA*

Abstract

Animals prefer a small, immediate reward over a larger delayed reward (time discounting). Lesions of the orbitofrontal cortex (OFC) can either increase or decrease the breakpoint at which animals abandon the large delayed reward for the more immediate reward as the delay becomes longer. Here we argue that the varied effects of OFC lesions on delayed discounting reflect two different patterns of activity in OFC; one that bridges the gap between a response and an outcome and another that discounts delayed reward. These signals appear to reflect the spatial location of the reward and/or the action taken to obtain it, and are encoded independently from representations of absolute value. We suggest a dual role for output from OFC in both discounting delayed reward, while at the same time supporting new learning for them.

Keywords

reward; orbitofrontal cortex; delay; time discounting; value

INTRODUCTION

Should I stay or should I go? Every year, at the annual Society For Neuroscience (SfN) meeting, we ponder this question while waiting in an endless line for a cup of coffee. Late to the conference and in desperate need of caffeine, we wonder: is it really worth the wait (Fig. 1A)?

Starbucks has a lot riding on the answer to this question; but beyond that, it also addresses a fundamental issue underlying how neural systems value different rewards that differ in how quickly they can be obtained. In the lab, the neural mechanisms underlying this aspect of decision making are studied in tasks that ask animals to choose between a small reward delivered immediately and large reward delivered after some delay. As the delay to the large reward becomes longer subjects usually discount the value of the large reward, biasing their choices toward the small, immediate reward. Interestingly, lesions of the orbitofrontal cortex (OFC) alter the breakpoint at which animals abandon the large delayed reward for the more immediate reward, effectively influencing whether they “stay” or “go.”

Address for correspondence: Matthew Roesch, Department of Anatomy and Neurobiology, University of Maryland School of Medicine, 20 Penn Street HSF-2, Rm S251 Baltimore, MD 21201. Voice: 410-706-8910; fax: 410-706-2512. mroes001@umaryland.edu.

In this article, we review these studies in the light of recent data collected from our lab examining neural correlates of time discounting in OFC. We will argue that the varied effects of OFC lesions, which sometimes increase and other times decrease this breakpoint, reflect two patterns of neural activity in OFC, one that maintains representations of reward across a delay and the other that discounts delayed rewards. We will show that these representations of the discounted reward are maintained independently of representations of absolute reward value. This is consistent with the finding that lesions to OFC disrupt delay discounting but often do not affect preference between differently sized rewards. The independent encoding of different aspects of reward value contradicts recent hypotheses that OFC neurons signal the value of outcomes in a kind of “common currency.” Finally, we will discuss findings that suggest that OFC might represent the spatial location and/or the action associated with delayed reward.

What Does Neural Activity Tell Us about the Role of OFC in Time Discounting?

Several studies have reported abnormal behavior in OFC-lesioned rats forced to choose between small, immediate and larger, delayed rewards. Some studies report that OFC lesions make animals more impulsive, that is, less likely to wait for a delayed reward.¹⁻³ These results suggest that OFC is critical for responding to rewards when they are delayed. Other studies report that OFC lesions make animals less impulsive; that is, more likely to wait for the delayed reward,^{3,4} suggesting that OFC is critical for discounting or devaluing the delayed reward. We have recently found two different patterns of neural activity in OFC, which appear to map on to these roles.

One pattern is evident in neurons that fire in anticipation of the delayed reward; activity in these neurons is similar to outcome-expectant activity seen in other settings.⁵⁻¹³ In our study, we trained animals to respond to one of two wells located under a central odor port (Fig. 1B–D). We then manipulated how long the animal had to wait to receive reward after responding. Figure 2 shows neuronal activity when the reward was delayed from 1–4 sec (gray). Figure 2A shows an OFC neuron that fires after the response into the fluid well and during the anticipation of the delayed reward. Figure 2B illustrates population activity of 27 OFC neurons sharing the same characteristic. When rewards were delivered after a short delay, activity rose after the response but quickly declined after reward delivery (black). However, when reward was delayed by 3 sec (gray) activity continued to rise until the delayed reward was delivered, resulting in higher levels of activity for rewards that were delayed compared to those that were delivered immediately.

These signals appear to maintain a representation of the imminent delivery of reward during the delay. Such a representation would facilitate the formation of associative representations in other brain regions. For example, we have recently shown that input from OFC is important for rapid changes in cue–outcome encoding in basolateral amygdala.¹⁴ This deficit may be due to the loss of these expectancy signals, normally generated in OFC.⁷ Similarly, the effects of pretraining OFC lesions on delayed discounting may reflect the absence of these expectancies when associations with immediate or delayed rewards must be learned. Rats lacking this signal during training would encode associations with the large, delayed reward more weakly than associations with the small, immediate reward. As a result, rats with pretraining lesions would exhibit apparently “impulsive” responding for the more strongly encoded small, immediate reward. This interpretation is consistent with reports that lesions of OFC before training cause impulsive responding.

Interestingly, experience with the large, delayed reward can reduce impulsive choice in OFC-lesioned rats.² In this experiment, a large, delayed reward was initially pitted against a small, immediate reward. OFC-lesioned rats were more impulsive than controls. However, after a period of training in which both rewards were experienced at equal delays, the OFC-lesioned

rats were able to perform normally when they were returned to a setting in which the small reward was no longer delayed. These results are consistent with the idea that OFC facilitates initial learning and that lesions can make rats more impulsive due to the loss of expectancy signals. Recovery of function under these circumstances may reflect the formation of more normal strength associations with the large reward, due to the additional experience in the symmetrically delayed variant of the task. Alternatively, associations with the small reward may be weakened in OFC-lesioned animals when the delay is imposed, for the same reason that associations with the large reward are weakened during initial training. These hypotheses might be distinguished by examining the strength of downstream encoding in areas such as basolateral amygdala and nucleus accumbens in lesioned rats.

However, OFC lesions can also result in lower levels of impulsivity, as if signals involved in discounting the delayed reward are lost.⁴ Interestingly the majority of OFC neurons recorded in well-trained rats in our study seemed to perform this function. Unlike the example in Figure 2A, activity in these neurons did not bridge the gap between the response and reward delivery, but instead declined as the delay to reward increased (Fig. 2C). Figure 2D represents the average firing rate of 65 neurons that showed this characteristic. Under short delay conditions, these neurons fired in anticipation of and during delivery of reward (black), however, when the reward was delayed (gray), activity declined until reward was delivered. Remarkably, this activity was correlated with a decreased tendency of the rat to choose the long delay on future free choice trials (Fig. 3; chi-square, $P < 0.05$). Thus, activity in this population biased decision making toward immediate gratification.

The effects of posttraining OFC lesions on delayed discounting may reflect the absence of these discounting signals, since posttraining lesions would not affect formation of the associative representations, but would cause the rats to be unable to discount the value of the large reward during actual task performance. As a result, rats with posttraining lesions would be less impulsive.⁴ This interpretation is consistent with reports that lesions of OFC after training cause lower levels of impulsive response. Of course, it should be noted that one study has reported that posttraining lesions can induce elevated levels of impulsivity.² The amount of training in this study, though substantial, was still far less than that reported in Winstanley *et al.*⁴ suggesting the rats may have not formed strong associations prior to surgery.

In summary, the effects of OFC lesions on delayed discounting may reflect the nature of these two signals and their role at different stages during learning. Rats lesioned after being fully trained have normal associative representations but are simply unable to discount the value of the large reward due to the loss of OFC. This reflects the loss of the discounting signal illustrated in Figure 2 C–D, which was the predominant signal in our well-trained rats. In contrast, rats lesioned before any training may have an associative learning deficit that renders meaningless the loss of any delayed discounting function in OFC; these rats fail to normally encode the reward associations due to the loss of expectancy signals during learning. As a result, they exhibit apparently “impulsive” responding, selecting the small reward lever even at very short delays for the simple reason that the associative representations of value for this response are better encoded.

Do Neurons in OFC Signal Time-Discounted Rewards in a Common “Value” Currency?

The results described above indicate that OFC is critical for discounting the value of the delayed reward after learning.⁴ This is conceptually similar to the role OFC plays for devaluation in other settings.¹⁵⁻¹⁸ The proposal that OFC is performing the same function in discounting and devaluation is consistent with the idea that output from OFC provides a context-free representation of value. This hypothesis is supported by single-unit recording work^{6,8,10,11,13,19-23} and functional imaging studies,²⁴⁻²⁹ which show that activity in OFC seems to encode the value of different goals or outcomes in a common currency. For example, Tremblay

*et al.*⁸ showed that OFC neurons fire selectively after responding in anticipation of the different rewards and that this selective activity is influenced by the monkey's reward preference. Activity in anticipation of a particular reward differed according to whether the monkey valued it more or less than the other reward available within the current block of trials. It has been proposed that this signal integrates available information that impacts this judgment, providing a context-free representation of a thing's value.^{25,26}

If this hypothesis is correct, then neural activity that encodes the delay to reward should also be influenced by changes in reward magnitude, either at a single-unit or population level. Such covariance has been reported when delay and reward size are manipulated at the same time.^{23,30} However, in the study described above, we found that when delay and reward size were manipulated across different blocks of trials, analogous to the manipulations of reward preference made by Schultz and colleagues,⁸ OFC neurons maintained dissociable representations of the value of delayed and differently sized rewards. Thus neurons that fired more (or less) for immediate reward did not fire more (or less) for a larger reward when the delay was held constant. This is illustrated for the 65 neurons that fired more for rewards delivered immediately (Fig. 4C), as well as for those ($n = 27$) that fired more strongly in anticipation of the delayed reward (Fig. 4A). Neither population showed any relation to reward size (Fig. 4B, D).

As a result, expectancy signals in rat OFC in our study did not track relative reward preference. This is evident in the activity elicited by the small immediate reward in the trial blocks that differed by delay versus magnitude of reward. Neurons encoding relative value in a common currency should have responded less for the small immediate reward when it was pitted against the large reward at the same delay (i.e., when it was nonpreferred), but more for this same small reward when it was pitted against a delayed reward of equal size (i.e., when it was preferred). Yet the relative value of the small reward was not reflected in the activity of either population (Fig. 2, black vs. Fig. 4, gray; t -test, $P > 0.7380$) or in the counts of single neurons (chi-square; $P > 0.31$).

The fact that we were able to dissociate the effects of reward and delay on single-unit activity in OFC indicates that encoding of these different types of value information may involve different neural processes. This dissociation is perhaps not surprising considering recent behavioral data that support the view that learning about sensory and temporal features of stimuli involves different underlying systems (see Delamater's article, this volume). In addition, many mathematical models of value typically treat size and delay as separate variables in their equations.³ However, our data indicate that, despite several reports to the contrary,^{8, 22,23} OFC neurons do not always provide a generic value signal. Our ability to detect this difference may reflect a species difference; however, a more interesting explanation is that the difference may reflect the amount of training typically required in primate studies, which is usually much greater than the training given to a rat. It is possible that with extended training, OFC neurons become optimized to provide these generic value representations. This would have interesting implications as it would suggest that OFC might do a good job integrating commonly encountered variables into a common currency and would do less well integrating variables that are unique or rarely encountered into these calculations.

Interestingly, on a population level, there was relatively little impact of reward size on activity in OFC, even though the rats responded similarly to size and delay manipulations (Fig. 5). Thus, OFC may have a particularly fundamental role in discounting delayed rewards that is not necessary for encoding the absolute value of a reward. This is supported by the finding that lesions to corticostriatal structures disrupt delay discounting but typically do not alter size preference.^{1,4,31}

Do Neurons in OFC Signal Time-Discounted Rewards Dependent on the Response Required to Obtain Them?

So how is OFC representing responses that lead to immediate versus delayed rewards? Until recently, the involvement of OFC in encoding the action taken to receive reward has been largely neglected. This is in large part due to the finding that task-related activity in primate OFC is generally not dependent on the direction of the motor response, suggesting that OFC encodes the value of rewards independent of the actions required to obtain them. However, in these studies, response direction is typically not a predictor of reward; instead monkeys are generally highly trained to associated reward qualities with the visual properties of conditioned stimuli.^{5,8-11,13,22,23}

In contrast, recent rodent work has explicitly paired reward with direction. These studies have found OFC neurons to be directionally selective, suggesting that OFC is involved in monitoring either the spatial goal or the action taken to achieve that goal.³²⁻³⁴ For example, in a study by Feierstein *et al.*,³² rats made responses to either a left or right well based on odor-direction contingencies. Remarkably, nearly half of the neurons recorded in OFC were directionally selective, firing more for one direction but not the other across multiple task epochs. They also found that neurons sensitive to the outcome (rewarded or not rewarded) were modulated by the side the animal had gone to, integrating directional and outcome information that may be used to guide future decisions. Similarly, in our discounting study, we found that discounting signals were directionally specific (Fig. 2; solid vs. dashed). Moreover, this directional signal did indeed impact choice behavior (Fig. 3).

The fact that signals in OFC are directional is important because most discounting tasks typically increase delays for only one response. Indeed, OFC does not appear to be necessary when delays are increased for both responses at the same time.² The presence of correlates reflecting the value of different directional responses in OFC is interesting in light of recent data from Balleine and colleagues showing that OFC lesions do not affect changes in instrumental responding after devaluation.³⁵ Such behavioral changes are presumed to require action–outcome associations. A lack of effect of OFC lesions on instrumental devaluation suggests that directional correlates, such as those demonstrated by us or the Mainen group,³² are either not critical to behavior or are not encoding action–outcome contingencies. One alternative is that they may encode associations between sensory features of the action and the outcome. In the case of our report, such features might include the spatial position of the well, for example. This is supported by connections with areas that carry spatial information³⁶ and by a number of lesion studies that implicate OFC in spatial tasks.^{37,38} Future studies that minimize these unique sensory aspects of the response would help elucidate the nature of this representation and its role in impulsive behavior.

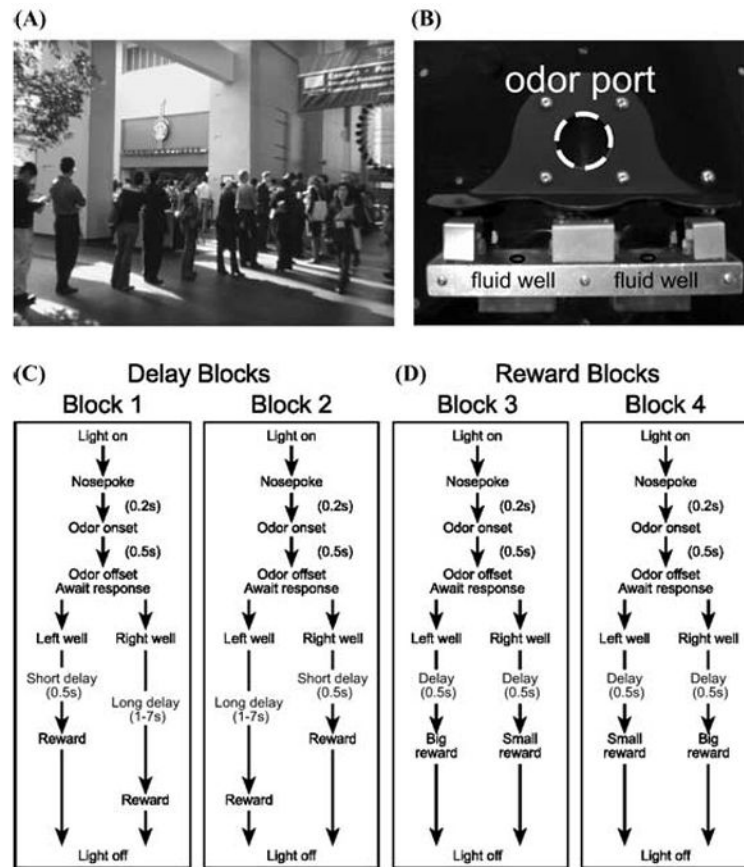
CONCLUSION

We conclude that the varied effects of OFC lesions on delayed discounting reflect two different patterns of activity in OFC; one that bridges the gap between a response and an outcome and another that discounts delayed reward. These signals appear to reflect the spatial location of the reward and/or the action taken to obtain it, and are encoded independently from representations of absolute value. This suggests a dual role for output from OFC in both discounting delayed reward, while at the same time supporting new learning for them. Output from OFC may impact reward representations in downstream areas shown to be involved in time discounting, such as basolateral amygdala and nucleus accumbens. Of course these predictions remain to be validated; the impact of delay on neuronal activity in these areas is unknown and remains critical to our understanding of how we decide to “stay” or “go” when rewards are delayed.

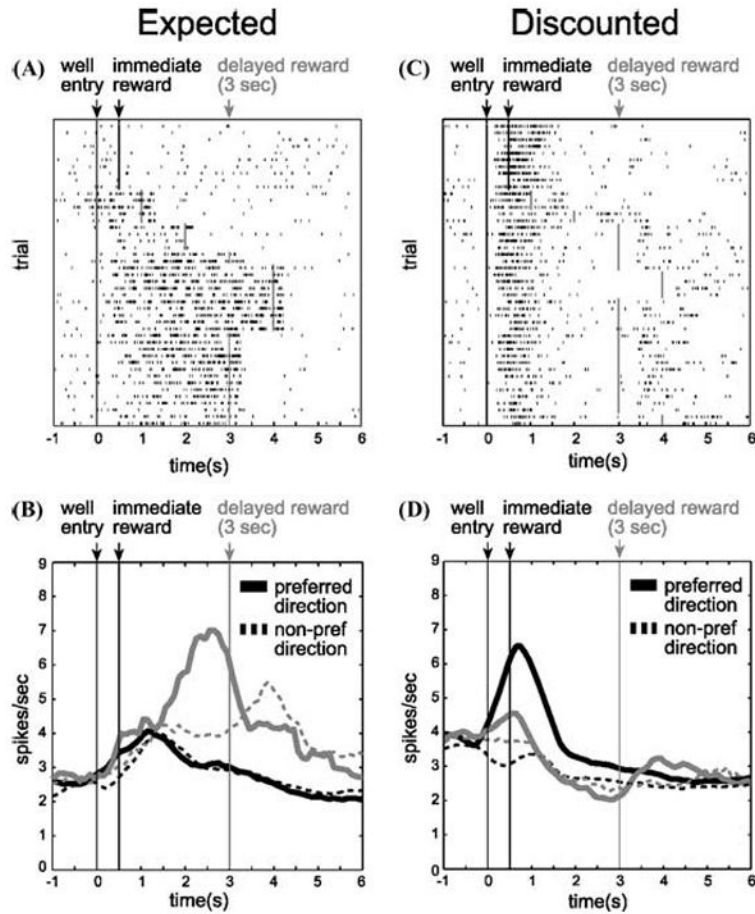
References

1. Mobini S, et al. Effects of lesions of the orbitofrontal cortex on sensitivity to delayed and probabilistic reinforcement. *Psychopharmacology (Berl)* 2002;160:290–298. [PubMed: 11889498]
2. Rudebeck PH, Walton ME, Smyth AN, et al. Separate neural pathways process different decision costs. *Nat Neurosci* 2006;9:1161–1168. [PubMed: 16921368]
3. Kheramin S, et al. Effects of quinolinic acid-induced lesions of the orbital prefrontal cortex on inter-temporal choice: a quantitative analysis. *Psychopharmacology (Berl)* 2002;165:9–17. [PubMed: 12474113]
4. Winstanley CA, Theobald DE, Cardinal RN, Robbins TW. Contrasting roles of basolateral amygdala and orbitofrontal cortex in impulsive choice. *J Neurosci* 2004;24:4718–4722. [PubMed: 15152031]
5. Rolls ET. The orbitofrontal cortex and reward. *Cereb Cortex* 2000;10:284–294. [PubMed: 10731223]
6. Schoenbaum G, Chiba AA, Gallagher M. Orbitofrontal cortex and basolateral amygdala encode expected outcomes during learning. *Nat Neurosci* 1998;1:155–159. [PubMed: 10195132]
7. Schoenbaum G, Roesch M. Orbitofrontal cortex, associative learning, and expectancies. *Neuron* 2005;47:633–666. [PubMed: 16129393]
8. Tremblay L, Schultz W. Relative reward preference in primate orbitofrontal cortex. *Nature* 1999;398:704–708. [PubMed: 10227292]
9. Tremblay L, Schultz W. Reward-related neuronal activity during gonogo task performance in primate orbitofrontal cortex. *J Neurophysiol* 2000;83:1864–1876. [PubMed: 10758098]
10. Wallis JD, Miller EK. Neuronal activity in primate dorsolateral and orbital prefrontal cortex during performance of a reward preference task. *Eur J Neurosci* 2003;18:2069–2081. [PubMed: 14622240]
11. Roesch MR, Olson CR. Neuronal activity related to reward value and motivation in primate frontal cortex. *Science* 2004;304:307–310. [PubMed: 15073380]
12. Hikosaka K, Watanabe M. Long- and short-range reward expectancy in the primate orbitofrontal cortex. *Eur J Neurosci* 2004;19:1046–1054. [PubMed: 15009152]
13. Hikosaka K, Watanabe M. Delay activity of orbital and lateral prefrontal neurons of the monkey varying with different rewards. *Cereb Cortex* 2000;10:263–271. [PubMed: 10731221]
14. Saddoris MP, Gallagher M, Schoenbaum G. Rapid associative encoding in basolateral amygdala depends on connections with orbitofrontal cortex. *Neuron* 2005;46:321–331. [PubMed: 15848809]
15. Schoenbaum G, Setlow B. Cocaine makes actions insensitive to outcomes but not extinction: implications for altered orbitofrontal-amygdalar function. *Cereb Cortex* 2005;15:1162–1169. [PubMed: 15563719]
16. Pickens CL, et al. Different roles for orbitofrontal cortex and basolateral amygdala in a reinforcer devaluation task. *J Neurosci* 2003;23:11078–11084. [PubMed: 14657165]
17. Izquierdo AD, Suda RK, Murray EA. Bilateral orbital prefrontal cortex lesions in rhesus monkeys disrupt choices guided by both reward value and reward contingency. *J Neurosci* 2004;24:7540–7548. [PubMed: 15329401]
18. Baxter MG, Parker A, Lindner CC, et al. Control of response selection by reinforcer value requires interaction of amygdala and orbital prefrontal cortex. *J Neurosci* 2000;20:4311–4319. [PubMed: 10818166]
19. Rolls ET. The orbitofrontal cortex. *Philos Trans R Soc Lond B Biol Sci* 1996;351:1433–1443. [PubMed: 8941955]discussion 1443–4
20. Critchley HD, Rolls ET. Olfactory neuronal responses in the primate orbitofrontal cortex: analysis in an olfactory discrimination task. *J Neurophysiol* 1996;75:1659–1672. [PubMed: 8727404]
21. Schoenbaum G, Chiba AA, Gallagher M. Neural encoding in orbitofrontal cortex and basolateral amygdala during olfactory discrimination learning. *J Neurosci* 1999;19:1876–1884. [PubMed: 10024371]
22. Padoa-Schioppa C, Assad JA. Neurons in the orbitofrontal cortex encode economic value. *Nature* 2006;441:223–226. [PubMed: 16633341]
23. Roesch MR, Olson CR. Neuronal activity in primate orbitofrontal cortex reflects the value of time. *J Neurophysiol* 2005;94:2457–2471. [PubMed: 15958600]

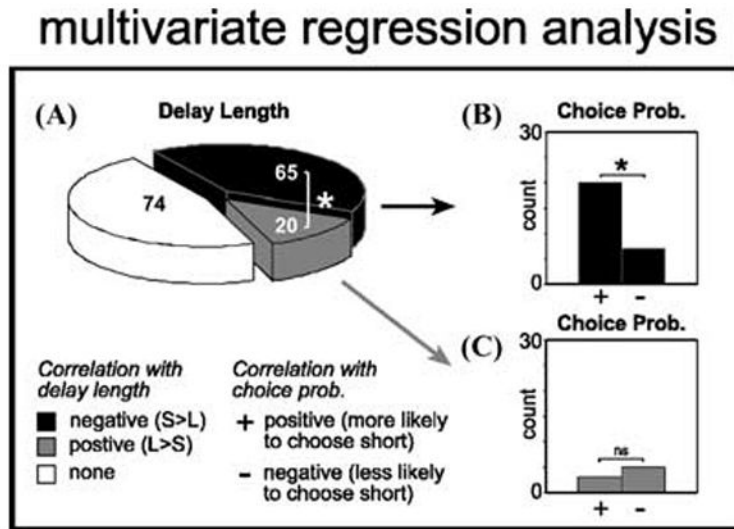
24. Gottfried JA, Doherty O’J, Dolan RJ. Encoding predictive reward value in human amygdala and orbitofrontal cortex. *Science* 2003;301:1104–1107. [PubMed: 12934011]
25. Kringelbach ML. The human orbitofrontal cortex: linking reward to hedonic experience. *Nat Rev Neurosci* 2005;6:691–702. [PubMed: 16136173]
26. Montague PR, Berns GS. Neural economics and the biological substrates of valuation. *Neuron* 2002;36:265–284. [PubMed: 12383781]
27. O’Doherty J, Critchley H, Deichmann R, Dolan RJ. Dissociating valence of outcome from behavioral control in human orbital and ventral prefrontal cortices. *J Neurosci* 2003;23:7931–7939. [PubMed: 12944524]
28. O’Doherty J, Kringelbach ML, Rolls ET, et al. Abstract reward and punishment representations in the human orbitofrontal cortex. *Nat Neurosci* 2001;4:95–102. [PubMed: 11135651]
29. Arana FS, et al. Dissociable contributions of the human amygdala and orbitofrontal cortex to incentive motivation and goal selection. *J Neurosci* 2003;23:9632–9638. [PubMed: 14573543]
30. Kalenscher T, et al. Single units in the pigeon brain integrate reward amount and time-to-reward in an impulsive choice task. *Curr Biol* 2005;15:594–602. [PubMed: 15823531]
31. Cardinal RN, Winstanley CA, Robbins TW, Everitt BJ. Limbic corticostriatal systems and delayed reinforcement. *Ann N Y Acad Sci* 2004;1021:33–50. [PubMed: 15251872]
32. Feierstein CE, Quirk MC, Uchida N, et al. Representation of spatial goals in rat orbitofrontal cortex. *Neuron* 2006;51:495–507. [PubMed: 16908414]
33. Roesch MR, Taylor AR, Schoenbaum G. Encoding of time-discounted rewards in orbitofrontal cortex is independent of value representation. *Neuron* 2006;51:509–520. [PubMed: 16908415]
34. Lipton PA, Alvarez P, Eichenbaum H. Crossmodal associative memory representations in rodent orbitofrontal cortex. *Neuron* 1999;22:349–359. [PubMed: 10069340]
35. Ostlund SB, Balleine BW. Lesions of the orbitofrontal cortex disrupt pavlovian, but not instrumental, outcome-encoding. *Soc Neurosci Abstracts* 2005;71.2
36. Reep RL, Corwin JV, King V. Neuronal connections of orbital cortex in rats: topography of cortical and thalamic afferents. *Exp Brain Res* 1996;111:215–232. [PubMed: 8891652]
37. Corwin JV, Fussinger M, Meyer RC, et al. Bilateral destruction of the ventrolateral orbital cortex produces allocentric but not egocentric spatial deficits in rats. *Behav Brain Res* 1994;61:79–86. [PubMed: 8031498]
38. Vafaei AA, Rashidy-pour A. Reversible lesion of the rat’s orbitofrontal cortex interferes with hippocampus-dependent spatial memory. *Behav Brain Res* 2004;149:61–68. [PubMed: 14739010]

**FIGURE 1.**


(A) Photograph taken at the annual neuroscience meeting in 2005 illustrates the all too familiar situation of waiting in an endless line for coffee. Late to the conference and in desperate need of caffeine you have to decide, do I go left and wait for coffee, or go right and see that last poster on my itinerary. (B) Apparatus used in our lab to study this type of decision making. (C–D). Choice task during which we varied (C) the delay preceding reward delivery and (D) the size of reward. Figure shows sequence of events in each trial in four blocks in which we manipulated the time to reward or the size of reward. Trials were signaled by illumination of the panel lights inside the box. When these lights were on, nosepoke into the odor port resulted in delivery of the odor cue to a small hemicylinder located behind this opening. One of three different odors was delivered to the port on each trial, in a pseudorandom order. At odor offset, the rat had 3 sec to make a response at one of the two fluid wells located below the port. One odor instructed the rat to go to the left to get reward, a second odor instructed the rat to go to the right to get reward, and a third odor indicated that the rat could obtain reward at either well. (C) One well was randomly designated as short and the other long at the start of the session (block 1). In the second block of trials these contingencies were switched (block 2). (D) In later blocks we held the delay preceding reward delivery constant while manipulating the size of the expected reward (adapted from Roesch *et al.*³³).

**FIGURE 2.**

(A) Single cell example of an expectancy neuron. Activity is plotted for the last 10 trials in a block in which reward was delivered in the cell's preferred direction after 500 msec (black) followed by trials in which the reward was delayed by 1 to 4 sec (gray). Each row represents a single trial, each tick mark represents a single action potential and the black/gray lines indicate when the reward was delivered. (B) Population histogram ($n = 27$) representing firing rate as a function of time during the trial for neurons that fired in anticipation of a reward delayed by 3 sec. Activity is aligned on well entry. Preferred direction refers to the spatial location that elicited the stronger response for each neuron. Black: short. Gray: long. Solid: preferred direction. Dashed: nonpreferred direction. (C) Single cell example of a neuron that discounts delayed rewards (gray), but fires strongly for rewards delivered immediately (black). (D) Population histogram representing firing rate as a function of time during the trial for neurons that fired more strongly after short delays during the reward epoch ($n = 65$). Conventions are the same as in A and B. "Preferred direction" refers to the spatial location of the well for which higher firing rate was elicited when other variables were held constant (i.e., reward size and delay length). Population activity is indexed to each neuron's preferred direction to average across cells (adapted from Roesch *et al.*³³).

**FIGURE 3.**

Dependency of firing rate during the reward epoch on delay length and future choice probability as revealed by multiple regression analysis. **(A)** Black and gray represent cases in which the correlation between firing rate and delay length was negative (stronger firing for short) or positive (stronger firing for long), respectively. **(B, C)**. Each bar represents the number of neurons in which the correlation between firing rate and future choice probability was positive (more likely to choose direction associated with short) or negative (less likely to choose direction associated with short) for those cells that also showed **(B)** a negative or **(C)** a positive correlation with delay length. * $P < 0.05$; chi-square. (adapted from Roesch *et al.* ³³).

**FIGURE 4.**

(A) Population histogram of same 27 neurons (shown in Fig. 2B) during trials when delay was held constant but reward size varied. Black: big. Gray: small. Solid: preferred direction. Dashed: nonpreferred direction. (B) Relation of firing dependent on delay length to firing dependent on reward size for those neurons that fired more strongly after long delays (shown in Fig. 2B). The delay index and reward index are computed on the basis of firing during the reward epoch. Delay index = $(S - L) / (S + L)$ where S and L represent firing rates on short- and long-delay trials, respectively. Reward index = $(B - S) / (B + S)$ where B and S represent firing rates on big- and small-reward trials, respectively. (C) Population histogram of same 65 neurons (shown in Fig. 2D) during trials when delay was held constant but reward size varied. Black: big. Gray: small. Solid: preferred direction. Dashed: nonpreferred direction. (D) Relation of firing dependent on delay length to firing dependent on reward size for those neurons that fired more strongly after short delays (shown in Fig. 2D). Conventions are the same as in A and B. (adapted from Roesch *et al.*³³).




FIGURE 5.

Impact of delay length (A–C) and reward size (D, E) on behavior. (A) Average choice rate, collapsed across direction, for all sessions for trials before and after the switch from short to long. Inset: The height of each bar indicates the percentage choice of short delay and long delay taken over all choice trials. (B, C) The height of each bar indicates the percentage correct (B) and reaction time (C) across all recording sessions in all rats on short-delay (black) and long-delay (gray) forced-choice trials. (D, E) Impact of reward size on the same behavior measures described in (A–C). Asterisks: t -test, $P < 0.05$. Error bars: standard errors. (adapted from Roesch *et al.*³³).