Neuronal Mechanisms in Prefrontal Cortex Underlying Adaptive Choice Behavior

JONATHAN D. WALLIS

University of California at Berkeley, Helen Wills Neuroscience Institute, Berkeley, California 94720, USA

ABSTRACT: **This chapter aims to address two questions relating to the role of the prefrontal cortex (PFC) in reward-guided choice behavior. First, do PFC neurons encode rewards per se, or are they encoding behavioral sequelae of reward? To address this, we recorded simultaneously from multiple PFC subregions, with the rationale that neuronal selectivity that directly encoded the reward outcome should occur before selectivity that reflected reward-related sequelae. Our results indicate that neurons in the orbitofrontal cortex (OFC) encode reward information before neurons in the dorsolateral PFC (DLPFC). Furthermore, whereas DLPFC neurons encoded both the upcoming response as well as the expected reward, OFC neurons encoded the reward alone. Our interpretation of these results is that the OFC encodes the reward and passes this information to the DLPFC, which uses it to determine the behavioral response. The second question is whether the encoding is specific to the reward outcome or reflective of a more abstract value signal that could facilitate decision making. We examined this by determining whether the PFC encodes other types of information relevant to decision making, such as probability of success and effort. We found that many PFC neurons encoded at least one of these variables, but neurons in the OFC and the medial PFC (MPFC) encoded combinations of the variables indicative of encoding an abstract value signal. This signal could provide decision making with flexibility and a capacity to deal with novelty, which are two of the hallmark features of prefrontal function. Future research will focus on delineating the differential contributions of the OFC and the MPFC to decision making.**

KEYWORDS: **monkey; neurophysiology; prefrontal; orbitofrontal; dorsolateral; reward; probability; decision making; choice**

Address for correspondence: Jonathan D. Wallis, University of California at Berkeley, Helen Wills Neuroscience Institute, 132 Barker Hall, Berkeley, CA, 94720.
wallis@berkeley.edu

Ann. N.Y. Acad. Sci. 1121: 447–460 (2007).
doi: 10.1196/annals.1401.009

INTRODUCTION

Making a bad choice can be serious. For example, the beta monkey that attacks the alpha male has to weigh the risk of injury or death against the biological need to procreate. How does the brain make such decisions? How does it ensure that it consistently selects the action most likely to realize the needs of the organism and to enhance the organism's survival prospects? The orbitofrontal cortex (OFC) is a key region in this regard, since damage to this area produces a relatively specific deficit in choice behavior. For example, consider the case of Elliott, a happily married young man in his 30s.[1,2] Elliott excelled in college and rose rapidly through the ranks of a home-building firm to become its chief accountant at the age of 32. Then, when Elliott was 35, doctors diagnosed him with a brain tumor. The operation to remove the tumor was successful, but the surgery left Elliot with bilateral damage to his OFC and the ventral portion of his medial prefrontal cortex (MPFC). However, neuropsychological tests of intelligence, memory, and language detected no evidence of brain damage. Even tests designed specifically to tax frontal lobe processes, such as working memory and rule switching, failed to reveal any deficits. Despite this, Elliot's life quickly spiraled out of control as he made a series of disastrous life decisions. He quit his job, lost a large sum of money to a scam artist, divorced his wife, lost contact with family and friends, and remarried a prostitute he had known for a month. His second marriage ended in divorce 6 months later, and he moved in with his parents. Thus, there is a paradox with the OFC: damage to this area leaves many of our cognitive abilities intact, yet it devastates our ability to make everyday decisions. In this chapter, we will focus on the underlying neuronal mechanisms that might help explain this paradox.

OFC NEURONS ENCODE EXPECTED REWARDS

The first neurophysiological studies of the OFC noted the frequency of neurons that showed responses to the delivery of juice rewards.[3,4] Subsequent studies showed that the neurons showed differential activity to two visual stimuli. One stimulus predicted the delivery of fruit juice, and the other predicted the delivery of saline.[4] Such neurons were not simply encoding the visual properties of the stimulus; when the reward contingencies were reversed, the neuronal selectivity also would reverse. Thus, the neurons appeared to be encoding the reward predicted by the stimulus and expected by the monkey. It was clear that these neuronal properties might be useful for decision making: encoding what reward to expect from a given action would allow the motor system to choose consistently the action that would lead to the largest reward.

However, the results of later studies countered the notion that these properties were unique to the OFC. For example, a series of studies demonstrated that neurons showing differential activity dependent on the expected reward

were also in the dorsolateral prefrontal cortex (DLPFC).[5–11] Particularly challenging was a study by Roesch and Olson, which examined the influence of expected reward magnitude on neurons throughout the frontal lobe.[12] Neurons showing a difference in firing rate depending on whether the subject expected a large or small reward were more prevalent in motor areas, such as the premotor cortex, than they were in the PFC. Similar neurons were also present in the posterior cortex, including the perirhinal cortex,[13] the parietal cortex,[14–16] and even the primary visual cortex.[17] We must be careful in interpreting these results, however. A neuron is not necessarily encoding a reward just because its firing rate correlates with some parameters of that reward. This is because many behavioral and cognitive measures also correlate with expected reward. For example, an animal's muscles often tense when it expects a large reward, and its behavior is quicker and more accurate.[12] An animal also pays more attention to cues that predict reward[18] and enters a state of higher autonomic arousal. Any of these processes may be driving neuronal firing rates.

How, then, do we determine whether OFC neurons are encoding the expected reward value or one of the correlates of reward? Our approach has been to compare the latency at which neurons encode expected rewards across various brain regions. Our rationale is that we must first determine that an animal expects a large reward before it can activate other cognitive processes, such as increased attention, arousal, and motor readiness. Thus, neurons that encode the expected reward will show differential activity dependent on the reward before neurons that encode cognitive processes that correlate with expected reward. In our first experiment, we recorded simultaneously from the DLPFC and the OFC to examine whether we could use this rationale to specify more precisely the contribution that both areas make to reward processing.

EXPERIMENT 1: COMPARISON OF REWARD ENCODING IN THE DLPFC AND THE OFC

In our first experiment we trained two monkeys to choose between different pictures associated with delivery of different amounts of fruit juice.[19] The subject would fixate on a central point on the screen, and two pictures would sequentially appear (one on the left and one on the right) separated by a delay (FIG. 1). The subject would then select one of the pictures by making a saccade to the location where that picture had appeared. Each picture was associated with the delivery of a specific amount of juice (0, 2, 4, or 8 drops). We used new pictures each day, and subjects learned by trial and error to maximize their reward by selecting pictures associated with larger juice amounts. Once subjects were consistently selecting the pictures associated with the largest reward (that is, they selected the largest reward on 27 out of the last 30 trials), we reversed the picture–reward contingencies. Thus, the picture that previously was associated with 8 drops of juice now was associated with 0 drops of juice,

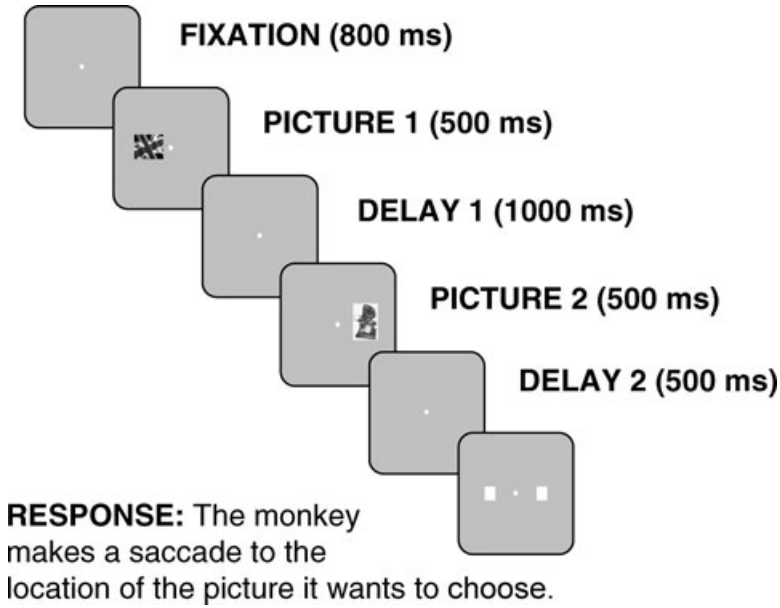

FIGURE 1. Illustration of the behavioral task that we used in Experiment 1.

the picture that previously was associated with 4 drops of juice now was associated with 2 drops of juice, and so on. This ensured that when a picture appeared on the screen, we could determine whether a neuron was encoding the reward that the picture predicted or was encoding the visual properties of the picture. For our present discussion, the most important neuronal activity was the one that occurred once the second picture appeared. At this point, the subject could predict what reward he would receive, as well as what motor response he would need to make in order to receive that reward.

We recorded neuronal activity simultaneously from multiple electrodes implanted in the DLPFC and the OFC. Recording simultaneously (as opposed to sequentially) from the two areas has the advantage that we are measuring the areas' neuronal activity during the exact same behavior, thereby controlling for subtle changes in behavior such as practice effects across recording sessions. We recorded the activity of 167 DLPFC neurons and 134 OFC neurons. FIGURE 2 illustrates two examples of OFC neurons that encoded the expected reward. FIGURE 2A and B show an example of an OFC neuron that encoded whether the subject expected to receive 4 drops of juice. The graphs show a higher firing rate on these trials compared to those in which the subject expected to receive either 2 or 8 drops. Its firing rate, however, was the same irrespective of whether the monkey made a left (FIG. 2A) or a right saccade (FIG. 2B). FIGURE 2C and D show an example an OFC neuron that encoded the expected reward in a parametric fashion. It showed a depression in its firing

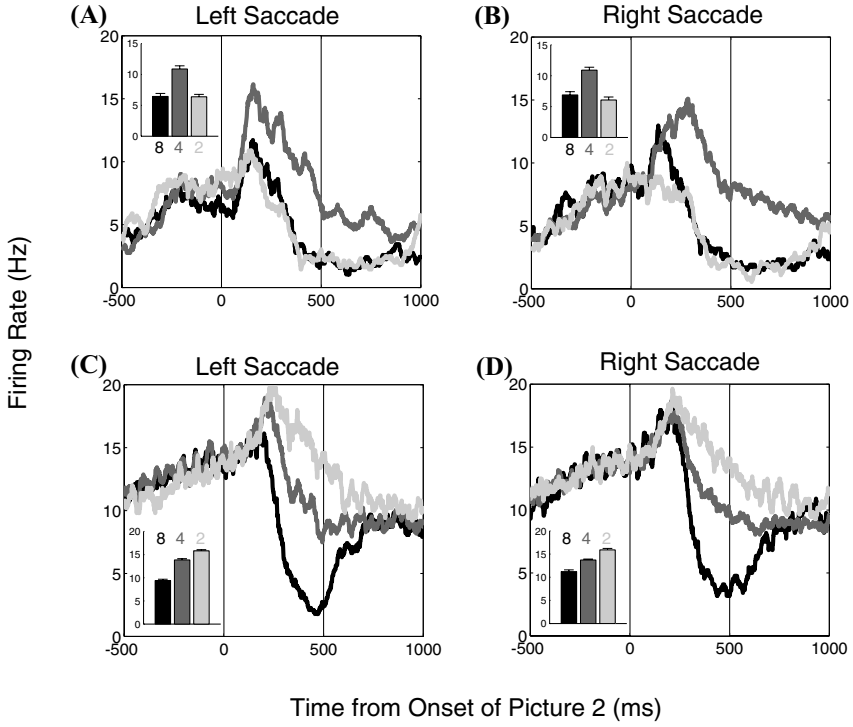

FIGURE 2. Spike density histograms from two OFC neurons indicating how the neuronal firing rate changed according to the expected payoff and the monkey's response (a left or right saccade). Inset bar graphs indicate the mean firing rate (± standard error) during the presentation of the reward-predictive cue (the first 500 ms). Black indicates that the cue predicted the delivery of 8 drops of juice, dark gray 4 drops, and light gray 2 drops. **(A, B)** OFC neuron showing a higher firing rate on trials in which the monkey expects to receive 4 drops of juice. **(C, D)** OFC neuron encoding the predicted reward in a parametric fashion. It showed a depression in its firing rate that was greatest for 8 drops of juice, less for 4 drops, and least for 2 drops. The upcoming saccade did not affect the firing rate of either neuron.

rate that was greatest for 8 drops of juice, less for 4 drops, and least for 2 drops. Again, the pattern of activity was independent of the direction of the saccade.

In contrast, DLPFC neurons tended to show complex responses that related to both the expected reward and the direction of the upcoming saccade. FIGURE 3 illustrates two representative DLPFC neurons. During the picture epoch, the neuron in FIGURE 3A and B discriminated between the different expected reward amounts when the monkey made a rightward saccade; it showed a high firing rate when 8 drops of juice was expected. (In contrast, during the subsequent delay period, the same neuron was reward-selective only when the

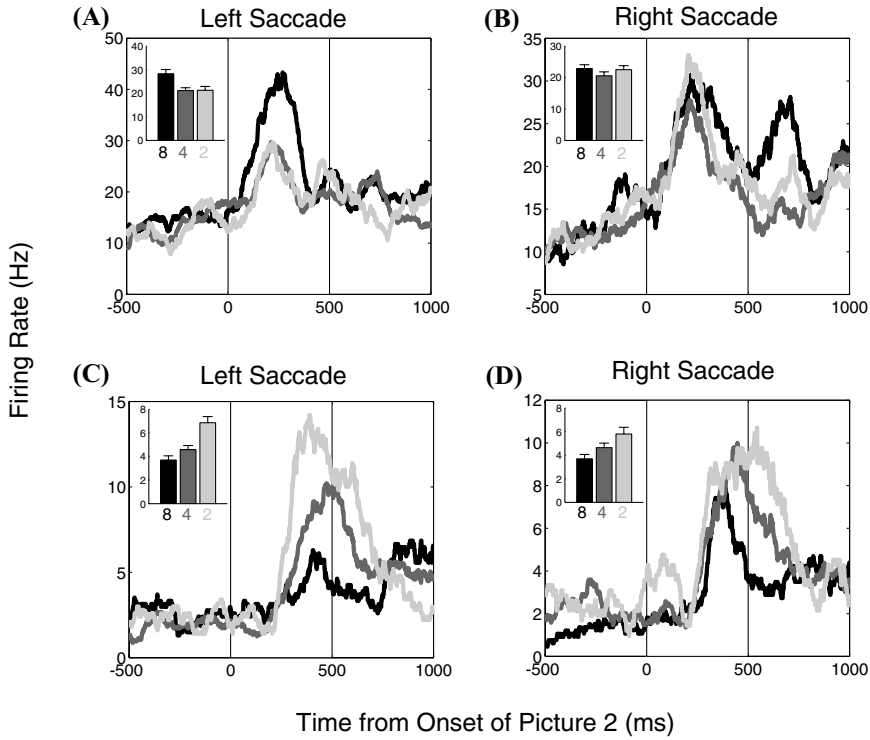

FIGURE 3. Spike density histograms from two DLPFC neurons, illustrated in the same manner as FIGURE 2. **(A, B)** During the time period that the picture was on the screen, the neuron discriminated between the different expected reward amounts, but only when the monkey made a rightward saccade. **(C, D)** The neuron encoded the reward in a parametric fashion, increasing its firing rate as the amount of expected reward increased. However, this effect was much greater when the subject was about to make a leftward saccade as opposed to a rightward saccade.

monkey made a leftward saccade.) The activity of the neuron in FIGURE 3C and D also was affected by both the upcoming saccade and the amount of juice that the subject expected. In this case, the neuron encoded the reward in a parametric fashion, increasing its firing rate as the amount of expected reward increased. However, this effect was much greater when the subject was about to make a leftward saccade as opposed to a rightward saccade.

These single neurons were representative of the properties of OFC and DLPFC neurons. To determine the proportion of these different types of neurons across the population, we performed a two-way ANOVA on the neuron's mean firing rate during the presentation of the second picture, using the factors of Reward (2, 4, or 8 drops of juice) and Saccade (leftward or rightward) (FIG. 4). In the OFC, 28% of the neurons showed a significant main

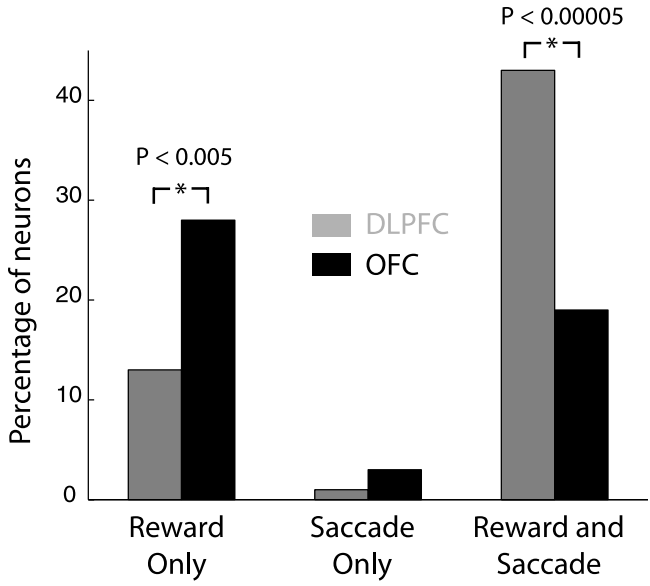

FIGURE 4. Bar chart illustrating the prevalence of neurons that encoded the expected payoff alone, the upcoming saccade alone, or an interaction between the payoff and the saccade. For every neuron that we recorded, we determined what it was encoding by performing a two-way ANOVA on the neuron's mean firing rate during the period that the picture was on the screen assessed at $P < 0.05$ (we saw similar results for the subsequent delay epoch). OFC neurons tended to encode the reward alone, while neurons in DLPFC encoded a combination of the reward and the upcoming saccade. Very few neurons in either area encoded the saccade alone.

effect of Reward (assessed at $P < 0.05$) with no main effect or interaction with Saccade, in comparison to 13% of DLPFC neurons ($\chi^2 = 9.8$, $P < 0.005$). In contrast, 43% of DLPFC neurons showed a significant Reward–Saccade interaction, compared to 19% of OFC neurons ($\chi^2 = 19$, $P < 0.00005$). Neurons in both areas encoded reward in a variety of ways. Some showed a parametric increase in firing rate as the expected reward size increased (27%). Others showed a parametric decrease (15%). Yet others encoded a specific reward (59%). These proportions were similar for both the OFC and the DLPFC.

A sliding receiver operating characteristic (ROC) analysis of the time-course of the selectivity revealed further differences in the encoding of reward between the two areas. Starting from 500 ms prior to the presentation of the second picture, we calculated an ROC value for a 200-ms time window. We then stepped this window forward in 10-ms increments until we had analyzed the rest of the trial. Briefly, the ROC analysis measured the degree of overlap between two response distributions. For each neuron, we defined the expected reward amount that yielded the highest firing rate as the preferred reward amount, and the expected reward amount that elicited the lowest mean firing rate as the

non-preferred reward amount. For trials in which the subject expected either the preferred or the non-preferred reward amount, we determined the total number of spikes that occurred in the 200-ms time window. This yielded two distributions of neuronal activity for trials in which the monkey expected either the preferred (P) or the non-preferred (N) reward. We then generated an ROC curve by taking each observed neuronal firing rate and plotting the proportion of P that exceeded the value of that observation against the proportion of N that exceeded the value of that observation. The area under this curve was then calculated. A value of 0.5 would indicate that the two distributions completely overlapped (since for each value of the neuron's firing rate the proportions of P and N exceeding that value are equal), and as such the neuron would not be selective. A value of 1.0, on the other hand, would indicate that the two distributions are completely separate (that is, every value drawn from N is exceeded by the entire distribution of P). In somewhat simpler terms, it is the probability that if I told you the firing rate of the neuron, you could predict which volume of juice the monkey expected to receive.

We used this analysis to compute the latency at which selectivity appeared. We defined this latency as the point at which the ROC curve exceeded 0.6. We chose the criterion as one that yielded a close approximation to the time at which we judged selectivity to appear from the spike density histograms. This measure did not differ between the two areas: 36% (60/167) of the DLPFC neurons reached criterion in a mean time of 467 ms, while 39% (52/134) of the OFC neurons reached criterion in a mean time of 426 ms (t-test $= 1.0$, $d.f. = 110$, $P > 0.1$). However, while selectivity for the reward tended to appear at about the same time in both areas, it then rose more rapidly and peaked earlier in the OFC than in the DLPFC (FIG. 5A). Therefore, for neurons that reached criterion, we calculated the value and time of the peak ROC value between the onset of the second picture and the start of the behavioral response. There was no difference between the two areas in the mean peak ROC value (DLPFC $= 0.654$, OFC $= 0.646$, t-test $= 0.89$, $d.f. = 110$, $P > 0.1$), but the peak was reached significantly earlier in the OFC than in the DLPFC. FIGURE 5B shows a distribution of the times at which each neuron reached its peak ROC value for the upcoming reward. On average, this occurred about 80 ms earlier in the OFC (510 ms after the onset of picture 2) than in the DLPFC (592 ms after the onset of picture 2; t-test $= 2.1$, $d.f. = 110$, $P < 0.05$).

In conclusion, we found neurons sensitive to the expected reward in both the DLPFC and the OFC. However, there was evidence for functional specialization. OFC neurons only encoded the expected reward, whereas DLPFC neurons encoded the upcoming saccade in addition to the expected reward. Further, OFC neurons encoded reward information 80 ms earlier than neurons in the DLPFC. The OFC is heavily and reciprocally connected with gustatory and olfactory cortices,[20,21] as well as the basolateral amygdala that might provide the OFC with information as to the value of the reward.[22,23] Thus, the OFC is conceivably the first prefrontal region that would receive information

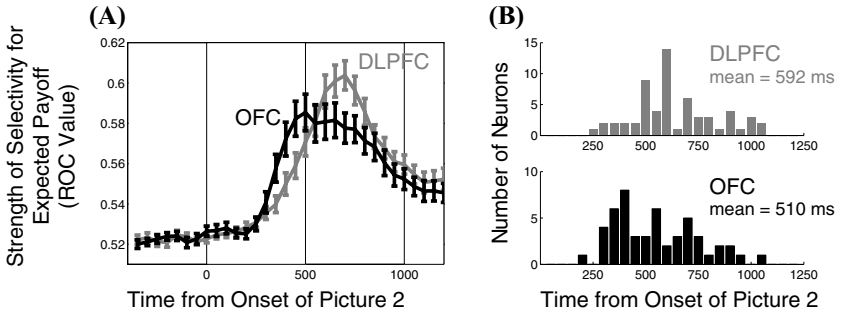

FIGURE 5. (A) Time-course of selectivity for the expected payoff (amount of reward) across the DLPFC (gray) and OFC (black) population of neurons. The thick line indicates the mean selectivity of the neurons, while the error bars indicate the standard error of the mean. Both populations began to encode the expected payoff at about the same time, but selectivity reached its peak value in the OFC before the DLPFC. The measure of selectivity is derived from the ROC of each neuron's firing rate. The ROC is the probability that an independent observer could correctly identify the payoff given the firing rate of the neuron. No selectivity equates to an ROC value of 0.5. (In practice it is slightly higher than this because we rectify the ROC value during its calculation. Small fluctuations due to noise push the value to about 0.52). Maximal selectivity equates to a value of 1.0. **(B)** Distribution of peak selectivity across the population of DLPFC and OFC neurons. The OFC population reached its peak selectivity approximately 80 ms before the DLPFC population (Wilcoxon's rank-sum test, $P < 0.05$).

about the value of the forthcoming juice reward. Our observation that reward value information peaks sooner in the OFC than in the DLPFC is consistent with that notion. In contrast, the timing of reward information in the DLPFC, along with these neurons' tendency to encode the upcoming response, suggest that this area may be where information about reward value converges with information about the subject's actions, thus allowing the subject to choose between the two different reward amounts. Thus, our hypothesis is that information about expected rewards enters the PFC through the OFC and then is relayed to the DLPFC. If this hypothesis is true, inactivation of the DLPFC should not affect reward information in the OFC, whereas inactivation of the OFC should attenuate reward information in the DLPFC. Future experiments will test this hypothesis.

EXPERIMENT 2: ENCODING OF OTHER IMPORTANT DECISION PARAMETERS

Thus far, we have seen that OFC neurons show differential activity depending upon the reward that the subject expects for a given choice. We have also seen that the timing of this activity is more consistent with the OFC neuronal response reflecting the encoding of the reward, rather than reflecting

a cognitive process that merely correlates with the reward. We have suggested that this neuronal response would make an important contribution to decision making by indicating to the motor system the action that would lead to the larger reward. However, decision making is more complex than simply always choosing the maximal reward. For example, a large reward may be obtainable, but if it is difficult to obtain it may be better to aim for a smaller, more readily obtainable reward.

Evolutionary biologists and economists have constructed detailed models of how we integrate different parameters to make effective decisions. These models emphasize the consideration of three basic parameters: the expected reward or payoff, the cost in terms of time and energy, and the probability of success.[24–26] Determining the value of a choice involves calculating the difference between the payoff and the cost and multiplying this by the probability of success. An obvious question, therefore, is to what extent OFC neurons also encode these other decision parameters. Do OFC neurons perform the calculations that are necessary for making an ideal choice, and is this their critical contribution to decision making? Recent results from neuroimaging studies suggest that this might be the case. The OFC (and sometimes the MPFC) is activated by manipulations of various decision variables—including probability,[27–29] payoff,[30–32] or the combination of these two variables—to create a set of integrated expected values.[33–37] These observations have led to the hypothesis that the OFC might integrate all variables relevant to making a decision to derive an abstract value signal, the so-called neuronal currency.[38]

A recent study in our own laboratory directly tested whether neurons in any of the major PFC regions were capable of responding to multiple parameters that underlie decisions. We trained monkeys to choose between pictures while we simultaneously recorded from the OFC, the MPFC, and the lateral PFC (LPFC).[39] Each picture was associated with a specific outcome. Some pictures were associated with a fixed amount of juice, but only on a certain proportion of trials (probability manipulation). Other pictures were associated with varying amounts of juice (payoff manipulation). Finally, some pictures were associated with a fixed amount of juice, but the subject had to earn the juice by pressing a lever a number of times (effort manipulation). About one-third of the PFC neurons responded parametrically to manipulations of just one of the decision parameters. These neurons occurred with equal prevalence in the three PFC areas from which we recorded. Furthermore, some neurons responded to a combination of two or more parameters. There was a progressive increase in the proportion of these neurons from the LPFC (16%) to the OFC (27%) to the MPFC (48%). Given these results, it is not surprising that the most severe decision-making deficits in humans occur after combined damage to the OFC and the MPFC.

We speculate that an important function of the OFC and the MPFC is to combine the multiple variables necessary to make a decision in order to

derive an abstract value signal. This simplifies the task of the motor system, which at any given instant should select the action with the highest value. This encoding scheme offers distinct computational advantages. When faced with two choices, A and B, one might imagine it would be simpler to compare them directly rather than going through an additional step of assigning them an abstract value. The problem with this is that as the number of available choices increases, the number of direct comparisons increases exponentially. Thus, choosing among A, B, and C would require three comparisons (AB, AC, and BC), while choosing among A, B, C, and D requires six comparisons (AB, AC, AD, BC, BD, and CD). The solution quickly suffers from combinatorial explosion as the number of choices increases. In contrast, valuing each choice along a common reference scale provides a linear solution to the problem.

An abstract representation provides important additional behavioral advantages, such as flexibility and a capacity to deal with novelty, both of which are hallmarks of prefrontal function. For example, suppose an animal encounters a new food type. In order to determine whether it is worth choosing relative to other potential food sources, the animal must determine the value of that food. If the animal relies on making direct comparisons, it can only determine the new food's relative worth by iteratively comparing it with all previously encountered foods. On the other hand, if the animal calculates an abstract value, it has to perform only a single calculation. By assigning the new food a value on the common reference scale, the animal knows the value of this foodstuff relative to all other foods. Second, it is often unclear how to compare directly very different outcomes. How does a monkey decide between grooming a conspecific and eating a banana? Valuing the alternatives along a common reference scale can help. For example, although I have never needed to value my car in terms of bananas, I can readily do so because I can assign the bananas and the car an abstract, monetary value.

Recent neuropsychological studies of decision making in patients with OFC and MPFC damage are consistent with this interpretation of our findings. Patients show unusual patterns of decision making when faced with complex choices that require the consideration and integration of multiple attributes.[40] For example, in choosing among different apartments, the patient might need to consider each apartment's size, neighborhood, and noise level. Some of these considerations might involve a trade-off between disparate variables, such as a large apartment in a so-so neighborhood or a small apartment in a good neighborhood, and thus would benefit from valuation along an abstract scale. Behavioral data suggested that controls attempted to make the choice that maximized as many of the attributes as possible, whereas patients followed a somewhat simpler strategy of assessing each apartment against some standard of acceptability, a strategy termed *satisficing.*

It is not just complex decisions that are impaired, however. Patients with OFC and MPFC damage also show erratic performance on a task that requires

preference judgments between stimuli presented two at a time, such as pictures of food, famous people, or even just colored swatches. Unlike controls, the patients showed erratic choices. For example, if they preferred A over B and B over C, they did not necessarily prefer A over C.[41] Thus, simple preference judgments also seem to benefit from the signals provided by the OFC and the MPFC.

CONCLUSION

Damage to the OFC produces a unique deficit. It impairs everyday decision making while leaving other cognitive capabilities intact. An extensive literature implicates the OFC in processing reward information, but in interpreting the area's neuronal responses, we must be careful to differentiate between responses that directly relate to the encoding of the reward's value and responses that only indirectly relate to the reward, via their encoding of cognitive and behavioral processes that covary with reward. Our findings from the first experiment suggest that the OFC is indeed encoding the reward's value, given the short latency of the neuronal reward-related responses. In contrast, the DLPFC appears to encode reward information as it relates to the guidance of behavioral responses. In addition, PFC neurons appear to encode other factors relevant to a decision, such as the effort required to obtain the reward and the probability of the reward's occurrence. Furthermore, we suggest that PFC neurons, particularly those in the OFC and the MPFC, are responsible for integrating the different decision variables to derive an abstract value signal. In turn, this signal would facilitate our capacity to make flexible and effective decisions in novel situations. Future research will aim to determine the precise contributions that the OFC and the MPFC make to decision making.

ACKNOWLEDGMENTS

Grants from NIDA R01-DA019028 and the Hellman Family Faculty Fund support our work. I would like to thank Earl Miller, in whose laboratory much of the work for the first experiment was completed. I would also like to thank Steven Kennerley for valuable conversations that went into the development of many of the ideas in this manuscript.

REFERENCES

1. ESLINGER, P.J. & A.R. DAMASIO. 1985. Severe disturbance of higher cognition after bilateral frontal lobe ablation: patient EVR. Neurology **35:** 1731–1741.
2. DAMASIO, A.R. 1994. Descartes' Error: Emotion, Reason, and the Human Brain. Putman. New York.

3. ROSENKILDE, C.E., R.H. BAUER & J.M. FUSTER. 1981. Single cell activity in ventral prefrontal cortex of behaving monkeys. Brain Res. **209:** 375–394.
4. THORPE, S.J., E.T. ROLLS & S. MADDISON. 1983. The orbitofrontal cortex: neuronal activity in the behaving monkey. Exp. Brain Res. **49:** 93–115.
5. HIKOSAKA, K. & M. WATANABE. 2000. Delay activity of orbital and lateral prefrontal neurons of the monkey varying with different rewards. Cereb. Cortex **10:** 263–271.
6. AMEMORI, K. & T. SAWAGUCHI. 2006. Contrasting effects of reward expectation on sensory and motor memories in primate prefrontal neurons. Cereb. Cortex **16:** 1002–1015.
7. KOBAYASHI, S. *et al.* 2002. Influence of reward expectation on visuospatial processing in macaque lateral prefrontal cortex. J. Neurophysiol. **87:** 1488–1498.
8. LEON, M.I. & M.N. SHADLEN. 1999. Effect of expected reward magnitude on the response of neurons in the dorsolateral prefrontal cortex of the macaque. Neuron **24:** 415–425.
9. WATANABE, M. 1990. Prefrontal unit activity during associative learning in the monkey. Exp. Brain Res. **80:** 296–309.
10. WATANABE, M. 1992. Frontal units of the monkey coding the associative significance of visual and auditory stimuli. Exp. Brain Res. **89:** 233–247.
11. WATANABE, M. 1996. Reward expectancy in primate prefrontal neurons. Nature **382:** 629–632.
12. ROESCH, M.R. & C.R. OLSON. 2003. Impact of expected reward on neuronal activity in prefrontal cortex, frontal and supplementary eye fields and premotor cortex. J. Neurophysiol. **90:** 1766–1789.
13. LIU, Z. & B.J. RICHMOND. 2000. Response differences in monkey TE and perirhinal cortex: stimulus association related to reward schedules. J. Neurophysiol. **83:** 1677–1692.
14. PLATT, M.L. & P.W. GLIMCHER. 1999. Neural correlates of decision variables in parietal cortex. Nature **400:** 233–238.
15. MUSALLAM, S. *et al.* 2004. Cognitive control signals for neural prosthetics. Science **305:** 258–262.
16. SUGRUE, L.P., G.S. CORRADO & W.T. NEWSOME. 2004. Matching behavior and the representation of value in the parietal cortex. Science. **304:** 1782–1787.
17. SHULER, M.G. & M.F. BEAR. 2006. Reward timing in the primary visual cortex. Science **311:** 1606–1609.
18. MAUNSELL, J.H. 2004. Neuronal representations of cognitive state: reward or attention? Trends Cogn. Sci. **8:** 261–265.
19. WALLIS, J.D. & E.K. MILLER. 2003. Neuronal activity in primate dorsolateral and orbital prefrontal cortex during performance of a reward preference task. Eur. J. Neurosci. **18:** 2069–2081.
20. MORECRAFT, R.J., C. GEULA & M.M. MESULAM. 1992. Cytoarchitecture and neural afferents of orbitofrontal cortex in the brain of the monkey. J. Comp. Neurol. **323:** 341–358.
21. CARMICHAEL, S.T. & J.L. PRICE. 1995. Sensory and premotor connections of the orbital and medial prefrontal cortex of macaque monkeys. J. Comp. Neurol. **363:** 642–664.
22. BAXTER, M.G. & E.A. MURRAY. 2002. The amygdala and reward. Nat. Rev. Neurosci. **3:** 563–573.
23. CARDINAL, R.N. *et al.* 2002. Emotion and motivation: the role of the amygdala, ventral striatum, and prefrontal cortex. Neurosci. Biobehav. Rev. **26:** 321–352.

24. KAHNEMAN, D. & A. TVERSKY. 2000. Choices, Values and Frames. Cambridge University Press. New York.
25. LOEWENSTEIN, G. & J. ELSTER. 1992. Choice Over Time. Russel Sage Foundation. New York.
26. STEPHENS, D.W. & J.R. KREBS. 1986. Foraging Theory. Princeton University Press. Princeton.
27. ABLER, B. *et al.* 2006. Prediction error as a linear function of reward probability is coded in human nucleus accumbens. Neuroimage **31:** 790–795.
28. CRITCHLEY, H.D., C.J. MATHIAS & R.J. DOLAN. 2001. Neural activity in the human brain relating to uncertainty and arousal during anticipation. Neuron **29:** 537–545.
29. HUETTEL, S.A. *et al.* 2006. Neural signatures of economic preferences for risk and ambiguity. Neuron. **49:** 765–775.
30. BREITER, H.C. *et al.* 2001. Functional imaging of neural responses to expectancy and experience of monetary gains and losses. Neuron **30:** 619–639.
31. ELLIOTT, R. *et al.* 2003. Differential response patterns in the striatum and orbitofrontal cortex to financial reward in humans: a parametric functional magnetic resonance imaging study. J. Neurosci. **23:** 303–307.
32. KNUTSON, B. *et al.* 2001. Anticipation of increasing monetary reward selectively recruits nucleus accumbens. J. Neurosci. **21:** RC159.
33. YACUBIAN, J. *et al.* 2006. Dissociable systems for gain- and loss-related value predictions and errors of prediction in the human brain. J. Neurosci. **26:** 9530–9537.
34. TOBLER, P. N. *et al.* 2007. Reward value coding distinct from risk attitude-related uncertainty coding in human reward systems. J. Neurophysiol. **97:** 1621–1632.
35. O'DOHERTY, J. *et al.* 2001. Abstract reward and punishment representations in the human orbitofrontal cortex. Nat Neurosci. **4:** 95–102.
36. KNUTSON, B. *et al.* 2005. Distributed neural representation of expected value. J Neurosci. **25:** 4806–4812.
37. DREHER, J.C., P. KOHN & K.F. BERMAN. 2006. Neural coding of distinct statistical properties of reward information in humans. Cereb. Cortex. **16:** 561–573.
38. MONTAGUE, P.R. & G.S. BERNS. 2002. Neural economics and the biological substrates of valuation. Neuron **36:** 265–284.
39. KENNERLEY, S.W., A.H. LARA & J.D. WALLIS. 2005. Prefrontal neurons encode an abstract representation of value. Society For Neuroscience. 194.16.
40. FELLOWS, L.K. 2006. Deciding how to decide: ventromedial frontal lobe damage affects information acquisition in multi-attribute decision making. Brain. **129:** 944–952.
41. FELLOWS, L.K. & M.J. FARAH. 2007. The role of ventromedial prefrontal cortex in decision making: judgment under uncertainty or judgment per se? Cereb. Cortex. **17:** 2669–2674.