Neuron
Article

Coding of Reward Risk by Orbitofrontal Neurons Is Mostly Distinct from Coding of Reward Value

Martin O'Neill[1,*] and Wolfram Schultz[1]
[1]Department of Physiology, Development, and Neuroscience, University of Cambridge, Downing Street, Cambridge CB2 3DY, UK
*Correspondence: mo293@cam.ac.uk
DOI 10.1016/j.neuron.2010.09.031

SUMMARY

Risky decision-making is altered in humans and animals with damage to the orbitofrontal cortex. However, the cellular function of the intact orbitofrontal cortex in processing information relevant for risky decisions is unknown. We recorded responses of single orbitofrontal neurons while monkeys viewed visual cues representing the key decision parameters, reward risk and value. Risk was defined as the mathematical variance of binary symmetric probability distributions of reward magnitudes; value was defined as nonrisky reward magnitude. Monkeys displayed graded behavioral preferences for risky outcomes, as they did for value. A population of orbitofrontal neurons showed a distinctive risk signal: their cues and reward responses covaried monotonically with the variance of the different reward distributions without monotonically coding reward value. Furthermore, a small but statistically significant fraction of risk responses also coded reward value. These risk signals may provide physiological correlates for the role of the orbitofrontal cortex in risk processing.

INTRODUCTION

Uncertainty is a ubiquitous component of our environment, such that humans and animals are regularly confronted with conditions that vary in degrees of uncertainty. It is therefore a fundamental requirement of our brain systems to accurately process uncertain information so that we may function appropriately, both in our mundane day-to-day activities and in more profound moments.

Many brain regions including the frontal cortex, basal ganglia, amygdala, parietal cortex, cingulate cortex, and insular cortex have been identified as key areas in processing information about uncertain rewards (Christopoulos et al., 2009; Fiorillo et al., 2003; Huettel et al., 2006; Knoch et al., 2006; Kuhnen and Knutson, 2005; Levy et al., 2010; McCoy and Platt, 2005; Preuschoff et al., 2006, 2008; Tobler et al., 2009; Xue et al., 2009). In particular, a large amount of research has focused on the orbitofrontal cortex. Damage to the ventromedial and orbitofrontal cortex leads to altered patterns of risky choice behavior (Bechara et al., 1994; Clark et al., 2008; Hsu et al., 2005; Mobini et al., 2002; Pais-Vieira et al., 2007; Rogers et al., 1999; Sanfey et al., 2003). In accordance with these findings, activations in the orbitofrontal cortex vary explicitly with risk (Critchley et al., 2001; Hsu et al., 2005; Kepecs et al., 2008; Tobler et al., 2007). In addition, neurons in the orbitofrontal cortex encode the value of available reward in a large variety of behavioral situations and in relation to the subjects' preferences (Hikosaka and Watanabe, 2000; Kennerley et al., 2009; Padoa-Schioppa and Assad, 2006; Peters and Büchel, 2009; Roesch and Olson, 2004; Thorpe et al., 1983; Tremblay and Schultz, 1999; Wallis and Miller, 2003). Thus, the orbitofrontal cortex appears to play an important role in processing both reward risk and reward value.

Although the orbitofrontal cortex is implicated in processing reward risk and value information, the capacity for individual neurons in the orbitofrontal cortex to explicitly encode risk information is unknown. We define risk as a form of uncertainty measured by the variance or standard deviation of known probability distributions of reward magnitudes (second central moment), rather than the occasionally used common sensical probability of losing. The purpose of this study was to test whether reward risk is encoded by neurons in the orbitofrontal cortex and, if so, to what extent such a risk signal would be distinct from the known value signal in this cortical area. Since risk is known to influence the subjective value of reward, a value signal might also covary with risk without constituting a genuine risk signal. Therefore, to identify a genuine risk signal, we found it essential to search for a stand-alone neuronal risk response that did not also encode value. To achieve this, we designed a simple task in which monkeys were presented with visual cues with distinct information about either the risk or the value of an upcoming juice reward. We recorded neuronal data under nonchoice conditions in which the monkeys had to perform an ocular saccade to the same location at which a single cue had been presented (left or right). Based on previous studies investigating reward processing in the orbitofrontal cortex, we predominantly targeted areas 11 and 13 (Kobayashi et al., 2010; Padoa-Schioppa and Assad, 2006; Tremblay and Schultz, 1999; Wallis and Miller, 2003). We found populations of orbitofrontal neurons that tracked the risk in predicted and received rewards in a monotonic fashion. True to the requirement of genuine risk coding, most of these neurons failed to also encode reward value.

RESULTS

Behavioral Preference for Risk and Value

This experiment varied separately the risk and value of liquid rewards. We used a formal definition of risk characterized by

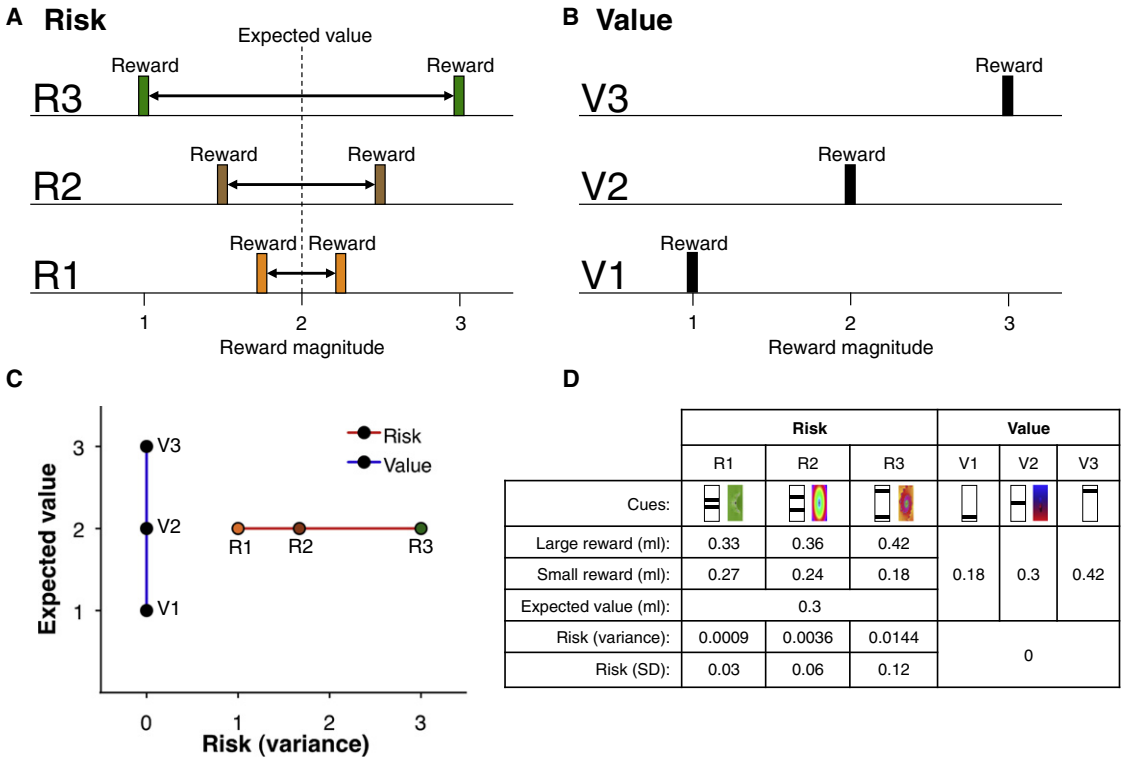

D

	Risk			Value		
	R1	R2	R3	V1	V2	V3
Cues:						
Large reward (ml):	0.33	0.36	0.42	0.18	0.3	0.42
Small reward (ml):	0.27	0.24	0.18			
Expected value (ml):	0.3					
Risk (variance):	0.0009	0.0036	0.0144	0		
Risk (SD):	0.03	0.06	0.12			

Figure 1. Independent Variation of Risk and Value

(A) Variations in reward risk. R1–R3 represent binary distributions of reward magnitude with equal probability (p = 0.5). R1–R3 differ in statistical variance (horizontal arrows), but not in expected value (dashed line).

(B) Variations in reward value. V1–V3 represent three reward magnitudes with fixed probabilities of p = 1.0.

(C) Graphical representation of the relationship between risk and value in (A) and (B).

(D) The trial types (R1–R3 and V1–V3), visual cues, and actual measures used in the experimental design. For the bar cues, the possible juice volumes to be delivered were represented by the elevation of the horizontal bars.

Rothschild and Stiglitz (1970), namely the mathematical variance of a known probability distribution. According to this definition the expected value (mean) is kept constant with different variances (the "mean-preserving spread"). We employed three levels of risk by using three different distributions with two equiprobable reward magnitudes (p = 0.5 each) and identical expected value. For example, as shown in Figure 1A, distribution R3 is the most risky distribution, whereas distribution R1 has the lowest risk. This definition of risk is consistent with those used in previous behavioral and neurobiological studies (Christopoulos et al., 2009; McCoy and Platt, 2005; Weber et al., 2004). We tested reward value without risk by varying the reward magnitude (Figure 1B). Figure 1C shows the relationship between the risk and value parameters. Figure 1D shows the actual measures of the reward distributions used in the experiment.

Two rhesus macaque monkeys performed a saccadic eye movement task for juice reward (Figure 2A). The animal touched a key and fixated on a spot in the center of a computer monitor while a visual cue was presented to the left or right of the spot. Our cues used the vertical position of horizontal bars to represent the magnitude of reward that would be delivered (the higher the bar, the greater the magnitude; Figure 1D). Risk conditions were indicated by cues containing bars at low and high vertical positions, representing the low and high equiprobable reward magnitudes, respectively. Cues with a single bar indicated that there was no risk and that the reward magnitude represented by the height of the bar would always be delivered (p = 1.0). To control for visual components, we also used a set of abstract fractal pictures for predicting outcome (Figure 1D).

We first quantified the monkeys' behavioral preference for the risky options in a choice task in which left or right eye movements indicated the monkeys' preferred option. In choice trials, monkeys chose between a safe option and a risky option (Figure 2B). The expected value of each risk option matched the expected value of the safe option (V2 in Figure 1D) to detect sensitivity to risk per se without confounding differences in expected value. Both monkeys preferred the riskier options to the safe option (Figure 2B; two-way ANOVA: main effect of risk, $F_{(2, 16)} = 48.74$, $p < 0.001$). The preference for risk was not significantly different between the two cue sets (main effect of cue set $F_{(1, 8)} = 0.18$, $p = 0.682$; cue set × risk interaction, $F_{(2, 16)} = 1.161$, $p = 0.338$). In addition, we measured reaction times to release the key during neuronal recording trials in which single risk or value cues were presented. Figure 2C shows that reaction times decreased as a function of increasing risk and value. Single linear regressions revealed a significant effect of

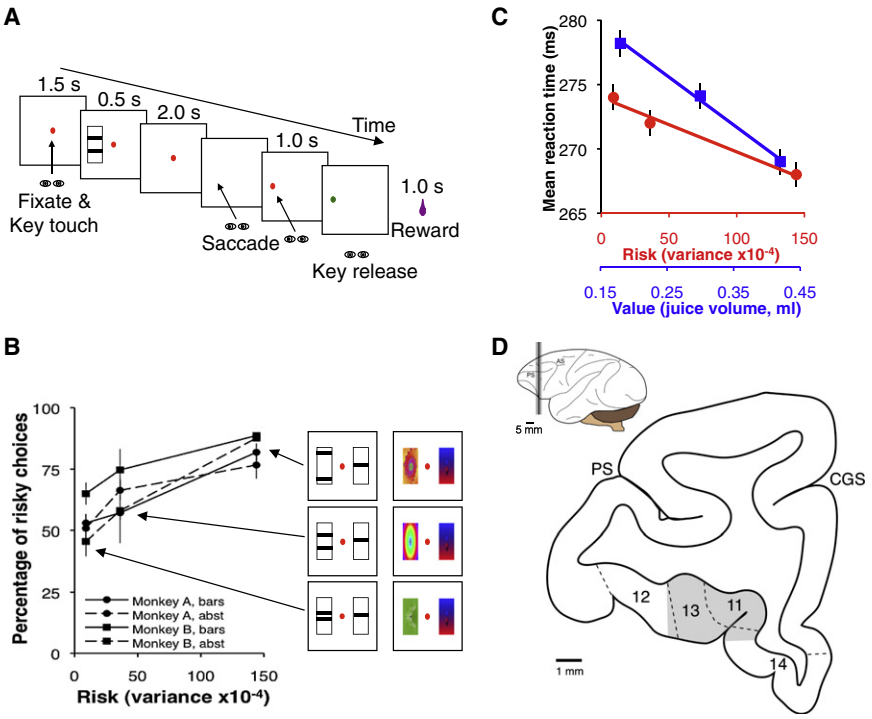

Figure 2. Behavioral Task, Measurements, and Neuronal Recording Area

(A) Behavioral task. Visual cues were presented on a monitor while monkeys fixated on a spot in the center of the screen and contacted a touch-sensitive key. For all neuronal recordings only one cue was displayed per trial, to the left or right of the fixation spot. After the fixation spot was extinguished the monkey was required to make a saccade to the side where the cue was displayed. After a successful saccade, a red dot appeared in this location for one second before turning green, indicating that the trial was complete and the key should be released to receive a juice reward.

(B) Behavioral preference for the risky options. In choice trials, two cues were presented simultaneously, on either side of the fixation spot. After the fixation spot was extinguished, the monkey was required to make a saccade to either side. Options were always between the medium value (safe) cue (V2 in Figure 1D) and one of the three risk cues. Each point in the graph shows the mean preference (±SEM), measured as percent of risky choices. In all three risk conditions both monkeys chose the risky option more than 50% on average and the preference for the risk option increased as a function of increasing risk (two-way ANOVA: main effect of risk, $F_{(2,\ 16)} = 48.74$, $p < 0.001$). The preference for risk was not significantly different between the two cue sets (main effect of cue set $F_{(1,\ 8)} = 0.18$, $p = 0.682$; cue set × risk interaction, $F_{(2,\ 16)} = 1.161$, $p = 0.338$).

(C) Behavioral discrimination of risk and value during neuronal recording trials. Data points are mean reaction times (±SEM) to release the key for reward from 114 blocks of trials. Reaction times decreased as a function of increasing risk ($p = 0.073$) and value ($p = 0.001$), and the slope for risk was not significantly different from the slope for value ($p > 0.1$; two-tailed t test).

(D) (Top) Sagittal view of the brain showing roughly the anterior-posterior extent of the region (thin vertical lines) where all neuronal data presented was recorded. (Bottom) Coronal view of a slice of the left hemisphere at the anterior-posterior region indicated by the thick vertical line in the upper image of the sagittal view. Shaded area denotes the targeted region of orbitofrontal cortex where all data presented was recorded from (±2.5 mm anterior-posterior). Numbers refer to architectonic areas. PS, principal sulcus; AS, arcuate sulcus; CGS, cingulate sulcus.

value ($p = 0.001$) and a lesser, not statistically significant, effect of risk ($p = 0.073$). The difference in regression slopes (β) between these regressions was nonsignificant ($p > 0.1$; two-tailed t test). These findings are consistent with previous evidence for risk-seeking behavior in macaques (Hayden and Platt, 2007; McCoy and Platt, 2005).

Neuronal Data Analysis and Database

We sampled electrophysiological activity from 722 single orbitofrontal neurons while the monkeys were presented with cues predicting different levels of reward risk and/or value. We recorded and saved the activity of neurons that appeared to respond to at least one task event during online inspection of several trials. This procedure resulted in a database from 262 orbitofrontal neurons, which we analyzed statistically in two consecutive steps. In the first step we defined task-related responses with the Wilcoxon test as significantly different neuronal activity during a given task epoch compared with that from a 1 s control period immediately preceding the central fixation spot. The task epochs analyzed were 0.1–0.6 s after cue appearance; 0.5 s immediately before the saccade, key release, and reward; 0.5 s immediately after the saccade and after key release; and four nonoverlapping, consecutive epochs of 0.5 s each after the reward. In the second step of the analysis, we assessed the monotonic neuronal coding of the two specific parameters of interest, namely risk and value of reward with single and multiple linear regressions on the task-related responses identified by the Wilcoxon test.

Risk Coding

The Wilcoxon test assessed all risky trial types together and identified a total of 1462 task-related responses in 262 neurons, with many neurons showing multiple task relationships. Single linear regressions using the mathematical variance as risk regressor ($9, 36, 144 \times 10^{-4}$) revealed that 160 of the 1462 task-related responses (11%) correlated significantly with risk ($p < 0.05$, two-tailed t test against 0 slope). Figure 3A shows four typical responses with positive correlation coefficients during four different task epochs. Figure 3B shows population averages of all responses with significant positive correlation coefficients during the same four task epochs. Figure 3C shows the numbers of responses with significant positive and negative correlation coefficients at each task epoch. Table 1 shows the numbers and percentages of significant risk coding responses

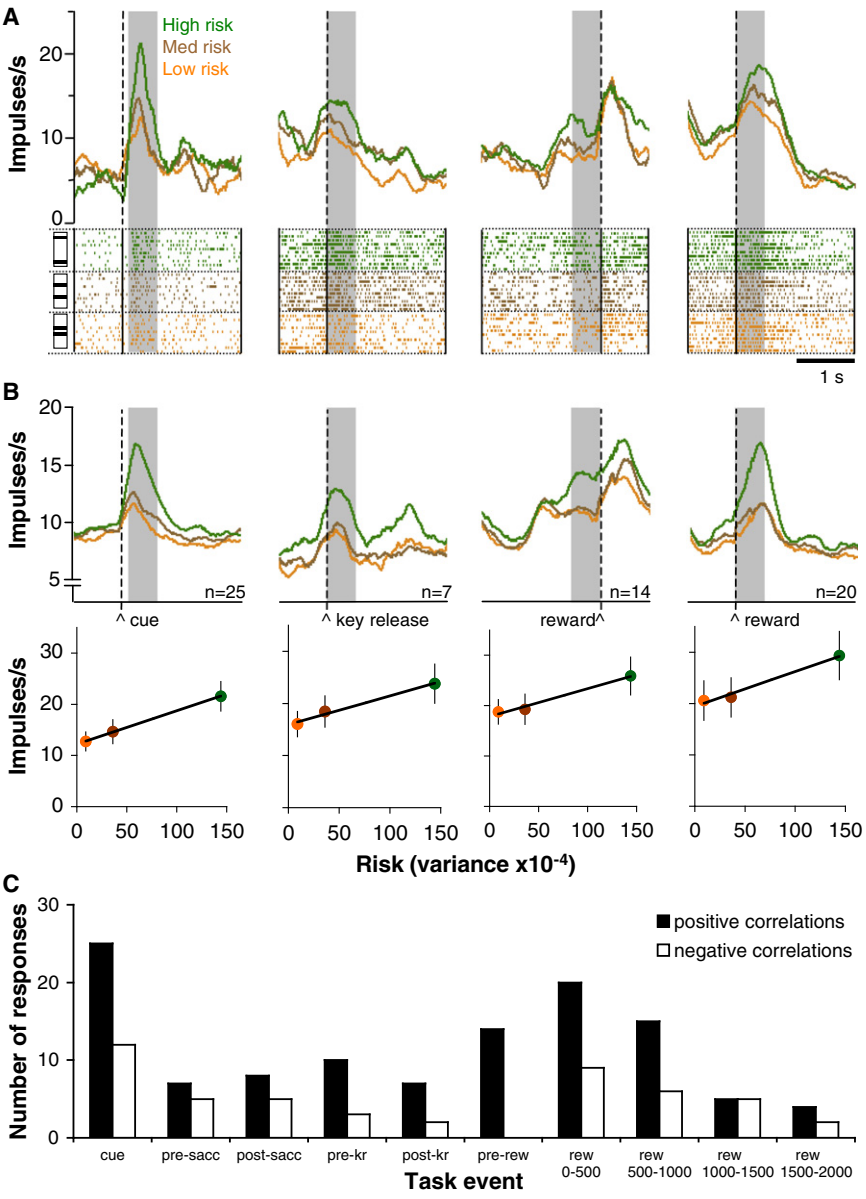

Figure 3. Orbitofrontal Neurons Code Risk

(A) Smoothed histograms showing examples of individual risk-sensitive neurons to the cue, key release, prereward, and postreward. For presentation purposes the raster plots are rearranged and blocked vertically by trial type (indicated by the cues beside the leftmost raster plots). The last ten trials for each trial type are displayed. The shaded area shows the time window used for analysis (correlation coefficients for all neurons, $p < 0.05$, two-tailed t test).

(B) Smoothed histograms showing population responses from all neurons with significant positive correlation coefficients for risk during the shaded period (top panels). The bottom panels show the mean firing rates (±SEM) during the shaded period in the upper panels. Orange = low risk; brown = medium risk; green = high risk.

(C) Numbers of risk-related responses with significant positive and negative correlation coefficients in all task epochs.

in each task epoch, amounting to 6%–18% of task-related responses in the different epochs. The 160 risk-related responses occurred in a total of 109 neurons. The majority of these neurons (79/109; 72%) showed a significant response to risk in only one task epoch, with some neurons showing significant responses in two (24/109; 22%) or more (6/109; 6%) task epochs. The 109 orbitofrontal neurons with risk-related responses were located in areas 11 (34 neurons), 12 (2 neurons), 13 (70 neurons), and 14 (3 neurons) (Figure 2D).

We also analyzed the 1462 task-related responses with standard deviation as risk regressor. Most risk responses displayed significant correlations for both variance and standard deviation. Figure 4 shows the close correlations between the r-squared values derived from the analysis with the two risk regressors (Pearson's $r = 0.97$, $p < 0.01$). Similarly small proportions of responses were significant for only one of the two risk measures (gray and white circles in Figure 4; $\chi^2 = 2$, $p = 0.2$; chi-square test), and varied closely around the statistical threshold ($p < 0.05$). Because the mean was constant across all trial types involving risk, the coefficient of variation (standard deviation/mean; Weber et al., 2004) as risk measure resulted in the same, albeit rescaled, results. Thus, these closely related risk measures produced very comparable results. For reasons of simplicity, we will present all data with variance as risk regressor.

Taking the square root of an independent variable steepens the slope of the relationship with the dependent variable. Indeed, because standard deviation is the square root of variance, the effect size measured by the slope (β) was larger for standard deviation than variance. The unsigned mean β values were 0.81 for standard deviation and 0.05 for variance.

It is possible that the neuronal responses to the cues reflected the visual properties of the cues rather than the risk per se (Wallis and Miller, 2003). Therefore, we tested 113 of the 262 task-related neurons with the three risky bar cues and the three risky abstract cues (Figure 1D). This sample included 19 neurons with cue responses that showed significant positive correlations to risk (they belonged to the 25 neurons with cue responses with positive slope shown in Figure 3B, left panels). Of the 19 neurons, 16 also showed significant positive risk correlation coefficients with the abstract risk cues (Figure 5); the remaining three neurons responded to the abstract cues in the same fashion as to the bar cues, but without reaching significance. Thus, the

Table 1. Significant Risk Responses in Each Task Epoch

Task epoch	Cue	Pre-sacc	Sacc	Pre-kr	Kr	Pre-rew	Rew1	Rew2	Rew3	Rew4	Total
Task-related	180	129	152	165	160	164	174	158	126	54	1462
Risk-sensitive	32	13	9	12	13	14	29	21	10	6	159
in percent	18	10	6	7	8	9	17	13	8	11	11
Slope (+/−)	25/12	10/3	7/2	7/5	8/5	14/0	20/9	15/6	5/5	4/2	115/44

Task-related responses were defined by the Wilcoxon test ($p < 0.05$), and risk sensitivity was defined by single linear regression ($p < 0.05$; t test against 0 slope). Task epochs were as follows: 0.1–0.6 s after cue onset (cue); 0.5 s presaccade (pre-sacc), 0.5 s prekey release (pre-kr); 0.5 s prereward (pre-rew); 0.5 s after saccade onset (sacc); 0.5 s after key release (kr); and four consecutive postreward epochs of 0.5 s each (rew1–rew4).

neuronal responses to the risky cues appeared to reflect the risk information conveyed by the cues rather than their visual properties.

Risk and Value Coding

Because the animals expressed behavioral preferences during both risk and value trials, it was necessary to determine the extent to which the neuronal responses reflected risk per se or some aspect of value associated with the risky outcomes. Therefore, we tested 493 of the sampled 722 orbitofrontal neurons with three value cues in addition to the three risky cues (Figure 1D). This design allowed us to collect neuronal data from six different, pseudorandomly alternating trial types in total, three of which varied in risk but had constant expected value, and three that varied in value with no risk (see Experimental Procedures for a description of the task design with separate risk and value trials). The Wilcoxon test identified 1083 task-related responses in 98 orbitofrontal neurons that were investigated in both risk and value trials; many neurons showed multiple task relationships. We analyzed the 1083 responses with two separate linear regressions on risk in the risk trials (variances of 9, 36, 144 × 10^{-4}) and on value in the value trials (magnitudes of 0.18, 0.30, 0.42 ml). The average number of trials per neuron was 13 for each of the three risk cues and 11 for each of the three value cues.

Of the 1083 task-related responses, 201 (19%) correlated significantly with risk, value, or both ($p < 0.05$, two-tailed t test against 0 slope). Of these 201 responses, 45 responses varied monotonically with risk only (Figure 6A), 138 responses varied monotonically with value only (Figure 6B), and 18 responses varied monotonically with both risk and value (Figure 6C), thus resulting in a total of 63 risk responses and 156 value responses (Tables 2 and 3). Of the 18 responses varying monotonically with both risk and value, 13 responses coded both variables with the same slope orientation and thus shared the same valence, whereas the remaining 5 responses showed opposite valence (green and gold sectors in Figure 6C, respectively). Figure 6D shows a scatter plot of the significant t values of the regression coefficients from risk and value responses.

To formally test whether the responses tended to code risk and value separately or in combination, we performed a chi-square test on a 2 × 2 contingency table of the significant and nonsignificant responses for risk and value (Table 3). The result suggested a significantly higher propensity for combined risk and value coding as opposed to separate coding of the two variables in the population of responses ($\chi^2 = 11$, $p = 0.001$).

The 201 risk- and value-related responses occurred in a total of 98 neurons. These neurons did not show differential distributions between orbitofrontal areas 11 and 13, in which most of these neurons were recorded ($n = 94$; $\chi^2 = 0.48$, $p = 0.49$; chi-square test). Separating the neurons according to their responses to risk only, to value only, and to risk and value in combination showed similar nondifferential anatomical distributions ($\chi^2 = 1.9$, $p = 0.38$). The very few neurons ($n = 4$) recorded in closely neighboring parts of area 14 were not included in the chi-square test. Taken together, the risk- and value-coding neurons described in this study were not differentially distributed in the regions of orbitofrontal cortex investigated (shaded area in Figure 2D).

We also compared the effect sizes for risk and value by using the slopes (β) of the regressions. Although risk and value ranges were chosen arbitrarily, some comparisons were possible within the ranges used. The effect size for risk was smaller compared with that for value when the mathematical variance

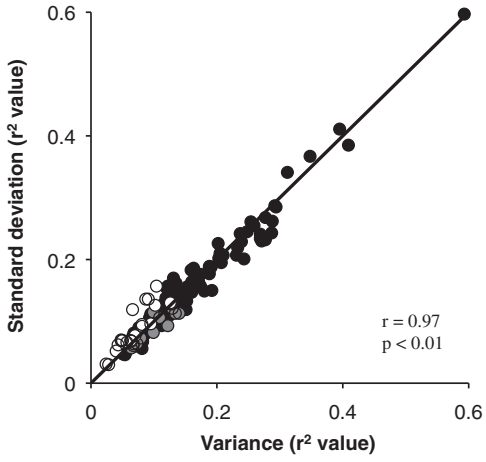

Figure 4. Mathematical Variance and Standard Deviation Were Both Considered as Measures of Risk

R-squared values from all neurons with significant correlation coefficients for variance and standard deviation from separate single linear regression analyses. Black circles represent neurons with significant r-squared values for both variance and standard deviation. White circles represent neurons with significant r-squared values for standard deviation and nonsignificant r-squared values for variance. Gray circles represent neurons with significant r-squared values for variance and nonsignificant r-squared values for standard deviation. The r-squared values were significantly related (Pearson's correlation). Line represents unity, along which both r-squared values would be equal.

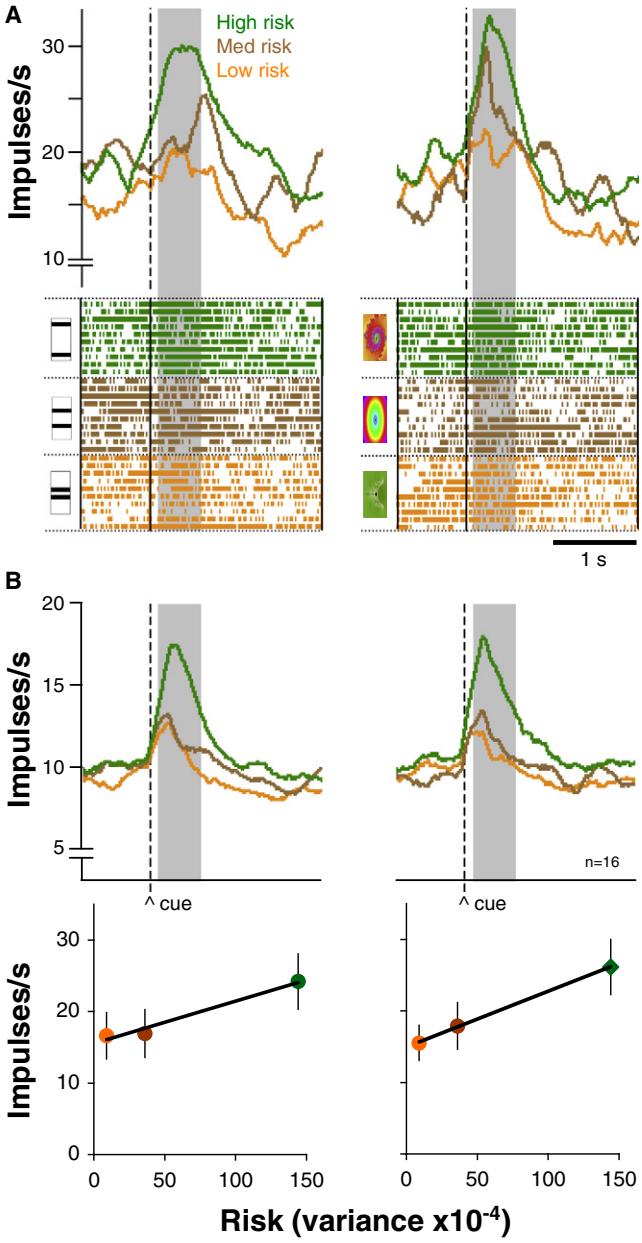

Figure 5. Cue Responses Reflect the Risk and Not the Visual Properties of the Cues

(A) Example of a single neuron with significant correlation coefficients for risk for both the risky bar ($p < 0.001$, two-tailed t test) and abstract cues ($p = 0.005$, two-tailed t test).

(B) Population responses of 16 neurons with significant correlation coefficients for both cue sets. (Conventions same as in Figure 3).

was used as risk measure (unsigned mean β values: $p < 0.001$; two-tailed t test), and greater than that for value with standard deviation (square root of variance) as risk measure (unsigned mean β values: $p < 0.001$). Thus, the effect sizes for the two risk measures straddled the effect size for value, suggesting similar effects of risk and value within the ranges of risk and value used.

Risk Coding after Outcome

The risk-sensitive neuronal responses to reward delivery shown in Figures 3A and 3B (right) and in Figure 6A were surprising because the risk had been resolved at that point. Furthermore, in the risky trials, the rewards occurred at the time of delivery in one of the two fixed magnitudes of each binary distribution, yet the risk-related responses appeared to covary with the risk, irrespective of the low and high reward magnitudes actually obtained. The fact that only a single magnitude of reward was delivered in risky trials afforded us with an additional approach for testing the coding of risk and value. Specifically, unlike the previous linear regression analysis where we compared separate single regressors on neuronal responses from separate trials, we were able to use multiple regressions on the same neuronal responses (to a single reward) in the same (risky) trials across the two regressors, risk and value (reward magnitude). The Wilcoxon test assessed the six risky trial types with their different reward outcomes separately and identified 641 task-related responses during one or more of the four postreward epochs in 231 orbitofrontal neurons, with many neurons showing multiple postreward responses. We analyzed the 641 task-related responses with multiple linear regressions using reward risk (variances of 9, 36, 144 × 10^{-4}) and value (reward magnitudes of 0.18, 0.24, 0.30, 0.33, 0.36, 0.42 ml) as regressors.

Of the 641 postreward responses, 42 responses (7%) showed significant partial correlation coefficients for risk, and 105 responses, for value (16%) ($p < 0.05$, two-tailed t test against 0 slope; Table 4). Figures 7A and 7B show the responses of neurons with positive slopes for reward magnitude and risk, respectively. Importantly, the responses in Figure 7B varied monotonically with the risk associated with each reward, independent of the reward magnitude (Figure 7C). For example, the juice volumes 0.18 and 0.42 ml were distinctly different in magnitude, but they were associated with the same amount of risk (see Figure 1A) and drew the same neuronal response. As before, for the chosen ranges of risk and value, the effect size for risk was smaller compared to that for value when variance was used as the risk measure (unsigned mean β values: $p < 0.001$; two-tailed t test) and greater than the effect size for value with standard deviation as the risk measure (unsigned mean β values: $p < 0.001$). Thus, plotting neuronal responses as a function of variance resulted in a flatter slope for risk compared with that for reward magnitude (Figures 7A and 7B).

However, as a possible alternative to risk coding, the pattern of activity observed in Figures 7B and 7C may represent nonmonotonic reward magnitude coding. In particular, the data from risky trials, when plotted across reward magnitude (Figure 7C), suggests quadratic reward magnitude coding. If this were the case, these neurons should code reward magnitude nonmonotonically also in nonrisk trials. To assess this possibility, we also tested 14 of the neurons presented in Figures 7B and 7C with the nonrisky value cues (Figure 7D). These 14 neurons showed 14 of the 32 risk responses with positive slope shown in Figure 7B. Fitting the neuronal responses in value trials with linear and quadratic functions revealed that 11 of these 14 risk neurons (79%) did not code reward magnitude (linear and quadratic least mean squares fits were both nonsignificant; F-test,

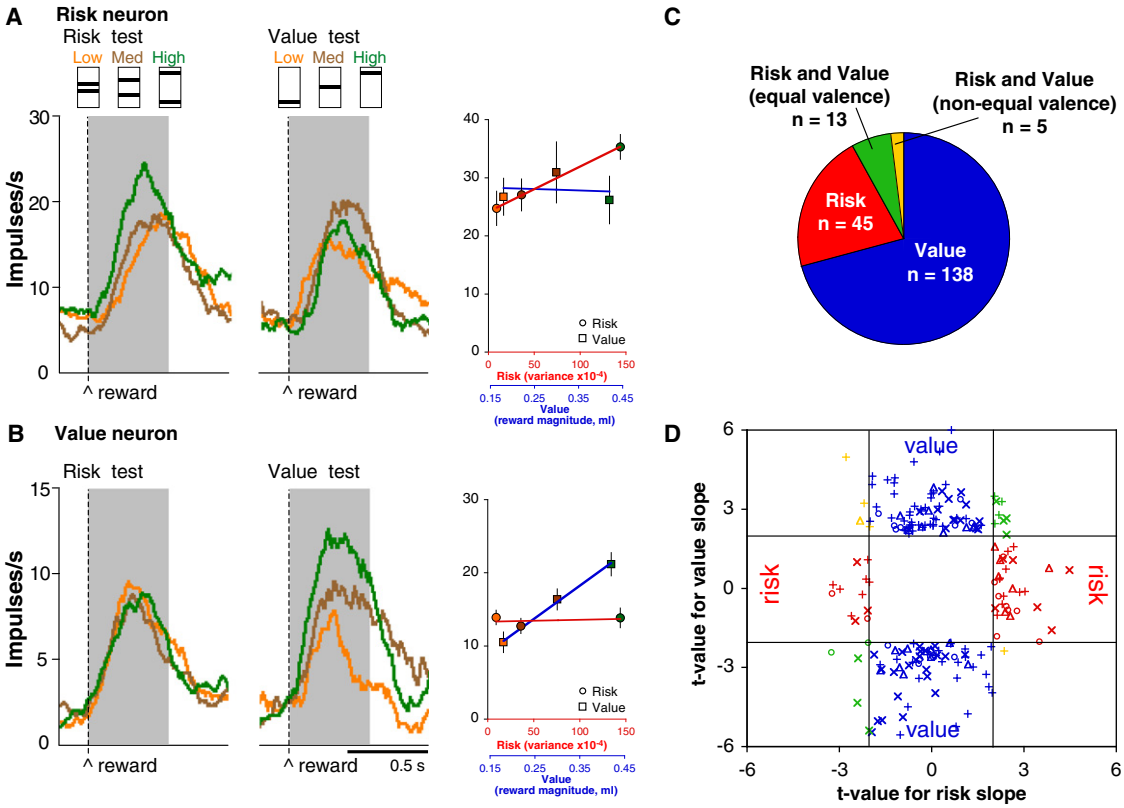

Figure 6. Risk and Value Coding

(A) Example of an orbitofrontal neuron that coded the risk associated with reward and not value. The left panel shows the neuron's firing rate during risk trials. The middle panel shows the firing rate during value trials. The right panel shows the average firing rate (±SEM) during the shaded periods. Regression lines were fitted to the risk (red) and value (blue) trials. Only the correlation coefficient for the risk trials was significant ($p < 0.005$, two-tailed t test).

(B) Example of an orbitofrontal neuron that codes the value and not the risk associated with reward. Conventions as in (A). Only the correlation coefficient for the value trials was significant ($p < 0.001$, two-tailed t test).

(C) Pie chart of the percentage of responses characterized by significant correlation coefficients for risk and/or value.

(D) Scatter plot of significant regression coefficients ($p < 0.05$, two-tailed t test) for risk (red), value (blue), or both (green, equal valence; gold, nonequal valence). The data points from responses that were nonsignificant for both variables are not shown for simplicity of display. Circles = saccade activity; triangles = key touch activity; plus sign = reward activity.

$p > 0.05$). Of the remaining three neurons, two showed significant correlation coefficients for both quadratic and linear functions (F-test, $p < 0.03$), and one showed a significant linear function only (F-test, $p = 0.04$). Therefore, the majority of these neurons appear to be coding reward risk rather than coding reward magnitude nonmonotonically. In addition, these data indicate that 21% (3/14) of reward risk coding responses potentially also code reward value, which is similar to the proportion of risk responses also coding reward value shown in Figure 6C (29%; 18/63; Table 2).

DISCUSSION

This study investigated the neuronal coding of reward risk, using the statistical variance of binary probability distributions of reward magnitudes as risk measure. We found that a group of orbitofrontal neurons coded information about reward risk in a positive or negative monotonic fashion across three risk levels. The risk coding occurred also with standard deviation as risk measure and was independent of the visual aspects of the risk-predicting cues. The majority of risk responses failed to vary monotonically with value, suggesting genuine coding of reward risk. However, a minority of risk-sensitive neurons also encoded reward value, and their propensity for value coding was significantly higher than that in the overall population of orbitofrontal neurons. Some orbitofrontal neurons also showed risk coding after the reward when the risk was resolved. Thus, orbitofrontal neurons appear to vary monotonically with the first two moments of probability distributions of reward magnitudes, namely expected value and variance (risk).

Neuronal Processing of Risk

Previous studies investigating the neural basis of risk processing in humans have used several versions of the Iowa Gambling task and variations in reward probability (Bach et al., 2009; Bechara et al., 1994; Clark et al., 2008; Critchley et al., 2001; Knoch et al., 2006; Kuhnen and Knutson, 2005; Levy et al., 2010; Peters

Table 2. Risk and Value Responses in Each Task Epoch

Task epoch	Cue	Pre-sacc	Sacc	Pre-kr	Kr	Pre-rew	Rew1	Rew2	Rew3	Rew4	Total
Task-related	126	94	98	115	109	103	126	123	100	89	1083
Risk-sensitive	16	7	5	4	4	3	11	7	4	2	63
in percent	13	7	5	3	4	3	9	6	4	2	6
Slope (+/−)	10/6	4/3	4/1	3/1	4/0	3/0	6/5	3/4	1/3	2/0	40/23
Value-sensitive	42	9	6	16	2	11	22	20	18	10	156
in %	33	10	6	14	2	11	17	16	18	11	14
Slope (+/−)	19/23	4/5	3/3	8/8	1/1	8/3	15/7	15/5	10/8	6/4	89/67
Both sensitive	7	1	1	1	0	1	1	4	1	1	18
in percent	6	1	1	1	0	1	1	3	1	1	2
Slope (+/−/op)	4/3/0	0/1/0	0/1/0	0/0/1	0	1/0/0	1/0/0	1/0/3	1/0/0	0/0/1	8/5/5

The numbers and percentages of 63 risk-sensitive and 156 value-sensitive responses include the 18 responses sensitive to both risk and value. Although 63 and 156 responses add up to 219, the total number of significant responses to both variables was 201 (= 156 + 63 – 18). op, opposite slopes for risk versus value. Task epochs were as in Table 1.

and Büchel, 2009; Xue et al., 2009). Other studies have assessed risk processing according to the definition of risk as mathematical variance or standard deviation of probability distributions used in economics and finance (Rothschild and Stiglitz, 1970). This definition provides a monotonic measure of risk, allowing a simpler and more straightforward identification of any biological response to the risk in the environment, as shown also by the use of linear regression models in our study. Using this approach, previous studies have identified risk processing in the human frontal cortex, parietal cortex, cingulate cortex, striatum, and amygdala (Christopoulos et al., 2009; Hsu et al., 2005; Huettel et al., 2006; Sanfey et al., 2003; Tobler et al., 2009). Certainly variations in probability can be expressed as nonmonotonic changes in risk, and this has been done to identify the involvement of frontal cortex, striatum, and dopamine neurons in risk processing (Fiorillo et al., 2003; Preuschoff et al., 2006; Tobler et al., 2007). However, task designs that manipulate variance while holding probability constant are more straightforward and have facilitated the identification of risk or utility signals in single neurons in the posterior cingulate cortex (McCoy and Platt, 2005) and of risk signals in our study in the orbitofrontal cortex.

The onset of differential risk-related activity we observed in the orbitofrontal cortex occurred as early as 100 ms after cue presentation. This latency appears to be shorter than the risk-related responses in dopamine neurons (Fiorillo et al., 2003) and also possibly that in cingulate neurons (McCoy and Platt, 2005). Specifically, the risk-related responses in dopamine neurons began relatively late after cue presentation and ramped up gradually until reward delivery. In the cingulate cortex, the risk or utility-related activity was most clearly observed after movements had been initiated. Therefore, the risk-sensitive orbitofrontal neurons observed in this study may constitute an early component of a system that processes risky information, perhaps before it is transmitted to other brain systems such as the dopamine neurons and cingulate cortex. This early response may possibly allow these orbitofrontal neurons to participate in detecting risk in decision situations involving uncertainty.

Table 3. Risk and Value Responses

	Value	Nonvalue	Total
Risk	18	45	63
Nonrisk	138	882	1020
Total	156	927	1083

Responses were grouped according to the statistical significance of the t values for the regression slopes for each variable. Responses prefixed with "non-" showed nonsignificant t values for the corresponding regressor.

Risk-Related Utility Coding

It is well known that risk influences the subjective value, or utility, of uncertain choices and determines behavioral preferences (Bernoulli, 1738). In common with our study, previous neurophysiological studies of risk processing did not directly address whether the risk activity observed was dissociable from a utility signal (Fiorillo et al., 2003; McCoy and Platt, 2005). A single-unit study in the orbitofrontal cortex that used a behavioral test for preferences provides a more direct assessment of utility coding (Padoa-Schioppa and Assad, 2006), but this study only employed certain, nonrisky outcomes. It would be possible to directly investigate a risk-related utility signal by using a similar design to that employed by Padoa-Schioppa and Assad (2006), but with risky options similar to those used in our task. Importantly, however, any such design would need to be able to distinguish between a subjective value signal for risk and a risk signal per se. For example, the risk signal in posterior cingulate cortex reported by McCoy and Platt (2005) was equivalent to the monkeys' behavioral preference and was therefore indistinguishable from a utility signal. Human studies have used behavioral preferences to identify risk signals that depend upon risk attitude (Christopoulos et al., 2009; Hsu et al., 2005; Huettel et al., 2006; Tobler et al., 2007). Furthermore, a recent study used this approach and provided evidence for combined coding of value and risk in the lateral prefrontal cortex (Tobler et al., 2009). Future single neuron studies of risk processing will benefit from distinguishing the subjective value of risky outcomes (utility) and the objective risk (mathematical variance).

Table 4. Risk and Magnitude Responses after Reward Delivery

Task epoch	Rew1	Rew2	Rew3	Rew4	Total
Reward-related	185	172	145	139	641
Risk-sensitive	19	12	7	4	42
in percent	10	7	5	3	7
Slope (+/−)	13/6	9/3	6/1	4/0	32/10
Magnitude-sensitive	27	30	26	22	105
in percent	15	17	18	16	16
Slope (+/−)	19/8	16/14	13/13	17/5	65/40

Responses to reward delivery in the four postreward periods were identified by the Wilcoxon test ($p < 0.05$). The risk and magnitude sensitivities were derived from subsequent multiple linear regressions in risk trials. Note that only a single reward magnitude occurred in risk trials at the time of reward delivery.

Although we did not test explicitly for differences in utility, the reaction time data seem to suggest a higher, albeit not significantly different, impact of value compared with risk within the ranges used with these two parameters (see Figure 2C). Based on this possibility, our neurons responding to both value and risk might have processed the utility associated with the different levels of risk. However, the value neurons that did not show risk coding may have been coding a utility signal that was undetectable with the risk cues because the utility variation due to risk might have been too small. Nonetheless, using standard deviation as risk regressor produced a larger effect of risk compared with value for the neuronal responses within the risk and value ranges used. If risk contributed to utility, these data suggest that the utility variation due to risk may not have been so small. Yet the responses of the value-only neurons failed to vary monotonically with the risk. Furthermore, the argument of potential utility coding would not apply to those risk-sensitive neurons that did not show value coding, as value coding would be necessary for utility coding.

It has been suggested that risk seeking might result from nonlinear weighting of rewards (Pratt, 1964). Although it is possible that the risk-seeking behavior of our monkeys was a result of an overweighting of the potential high rewards, the neuronal responses to risk alone are unlikely to be explained in this way because most of the neurons activated by the high-risk cue were not encoding value. Thus potential value coding does not seem to offer a general explanation for the observed risk coding in most of our risk-sensitive neurons.

Our finding that risk-coding neurons in the orbitofrontal cortex do not also frequently encode reward value monotonically is similar to previous findings indicating that populations of orbitofrontal neurons that are sensitive to probabilistic, costly, or delayed rewards are insensitive to absolute reward value (Kennerley and Wallis, 2009; Roesch et al., 2006). Therefore, reward outcome coding in orbitofrontal cortex appears to extend beyond a simple coding of reward value, incorporating additional information relating to outcomes, such as their riskiness, likelihood, cost, and temporal properties.

Neuronal Coding of Risk after Risk Is Resolved

The risk signal we observed following reward delivery might seem surprising given that risk has been resolved by this stage. However, these data are not unprecedented: previous studies have reported risk-dependent modulation of neuronal responses to the delivery of different reward magnitudes in the orbitofrontal

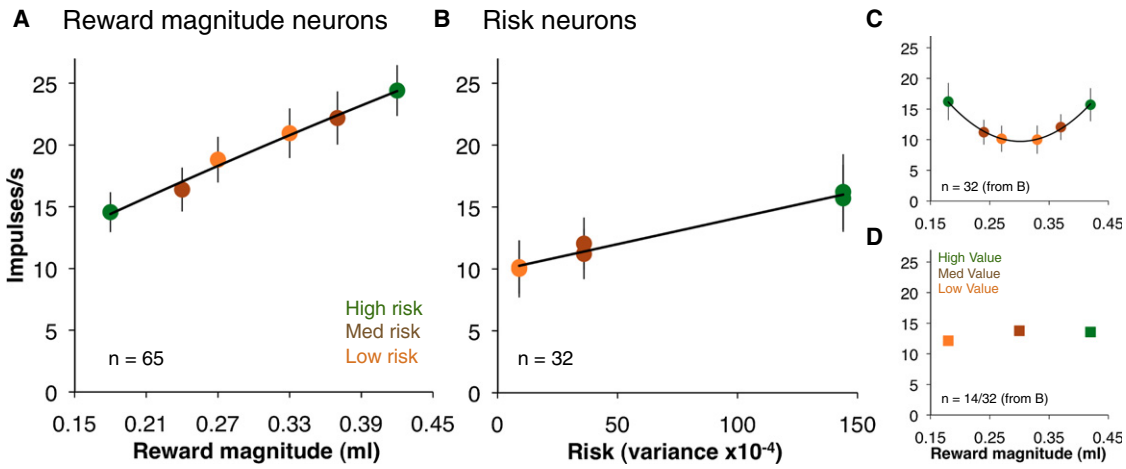

Figure 7. Orbitofrontal Neurons Code Risk Also after Risk Is Resolved

(A) Reward magnitude coding following reward in risk trials. Data points show average mean firing rates (±SEM) from 65 responses that had significant positive partial correlation coefficients for reward magnitude ($p < 0.05$, two-tailed t test).

(B) Risk coding following reward in risk trials. Postreward responses from risk trials are from a different population of neurons from (A), showing average mean firing rates (±SEM) from 32 responses that had significant positive partial correlation coefficients for risk ($p < 0.05$, two-tailed t test).

(C) Same data from (B) plotted across reward magnitude, showing the nonlinear, quadratic relationship with reward magnitude (curve shows best fit to quadratic function; F-test, $p = 0.036$).

(D) Risk responses did not code reward magnitude. Of the responses shown in (B), 14 were presented with both risk and nonrisky, value trials. Data points are the average mean firing rates (±SEM) of these responses from the separate nonrisky, value trials, showing the nondiscriminative responses to the fully predicted reward magnitudes.

cortex (Roitman and Roitman, 2010) and posterior cingulate cortex (Hayden et al., 2008). Although the latter study did not vary risk monotonically, the data are strikingly similar to ours: the study showed postreward responses to single outcomes that were greater with no risk than with risk. Similar neuronal responses at the time of outcome of uncertain decisions have also been observed in the anterior cingulate cortex (Quilodran et al., 2008). However, it should be noted that these risk-related responses following reward delivery may represent other forms of coding related to the reward, such as unsigned reward value prediction errors. Further investigations are warranted to clearly identify the relationship of postreward neuronal responses following the resolution of risk. Nonetheless, such alternative explanations do not apply to our data showing risk coding preceding the reward.

Functional Overview

Under conditions of risk, the distributions of possible outcomes are characterized by the mean outcome (expected value, the first moment of probability distributions) and the risk (mathematical variance, the second central moment). Interestingly, in financial economics, mean-variance theory (Levy and Markowitz, 1979) is based on this distinction between expected value and variance. Our data showing that some neurons in the orbitofrontal cortex encode these components separately would be compatible with this notion. However, there are other economic theories that provide alternative explanations for how reward distribution information is processed under conditions of uncertainty, such as expected utility theory (von Neumann and Morgenstern, 1944) and cumulative prospect theory (Tversky and Kahneman, 1992). Our data indicating a propensity of combined risk and value coding would not be incompatible with these two theories. Overall, these results might suggest physiological correlates for these theories but should not be taken to distinguish between their validity, because we have shown that risk and value are processed in different ways within the same brain structure.

Our findings offer an interpretation of the well-known dysfunctional decision making under conditions of uncertainty following damage to the orbitofrontal cortex in humans (Bechara et al., 1994; Clark et al., 2008; Hsu et al., 2005; Rogers et al., 1999; Sanfey et al., 2003). We have identified a monotonic risk signal in the orbitofrontal cortex that occurs in neurons that do not necessarily display a monotonic value signal. Thus, changes in decision making under uncertainty due to orbitofrontal damage may result from an inability to accurately process risk information. The absence of these neurons may be crucial for the deficient decision-making abilities of patients with orbitofrontal lesions. The deficit in decision-making under uncertainty in orbitofrontal patients may not represent a decision deficit of the highest cognitive order, but these patients may simply assess risks inappropriately because they do not have the risk signal propagated by the orbitofrontal neurons. Impaired risk processing would constitute a parsimonious explanation for the deficits in risky decision-making following damage to the orbitofrontal cortex and would provide a potential pathophysiological correlate for these striking and severely incapacitating behavioral deficits.

EXPERIMENTAL PROCEDURES

Subjects and Surgery

We used two adult male rhesus monkeys (*Macaca mulatta*), weighing 10–14 kg. The monkeys were implanted, under general anesthesia, with a head holder and stainless steel chamber on the skull to enable daily recordings of single neurons. All surgical and experimental procedures were performed under UK Home Office License in accordance with the requirements of the United Kingdom Animals (Scientific Procedures) Act 1986.

Behavioral Task

During training and testing, the monkeys were on a restricted water schedule 6 days of the week with 24 hr water ad libitum. The monkeys were trained to sit in a restraining chair in front of a computer monitor with the head fixed and perform a memory-guided saccade task. An aperture in the front of the chair provided access to a touch-sensitive key. To commence a trial, monkeys were required to fixate on a red spot in the center of the monitor and contact the key. After 1.5 s, cues were displayed, in a pseudorandomized order, to the left or right of the fixation spot for 0.5 s (Figure 2A). Fixation was maintained for a further 2 s before the center spot extinguished and this was the signal for the monkey to make a saccade to the left or right. A red fixation spot then appeared in this peripheral location and the monkey was required to fixate on this spot for 1 s before the spot turned green, which signified the end of the trial, and at which point monkeys could break fixation and were required to release contact with the key. Juice reward was delivered 1 s later. The beginning of the next trial was indicated by the appearance of the central fixation spot 3.5 s after the reward. Thus, the intertrial interval was 3.5 s, and the total cycle time (trial duration + intertrial interval) was 10.5 s. The task contained two categories of trials: imperative trials (as shown in Figure 2A), in which a single cue was shown at either side of the central fixation spot and the monkeys were required to make a saccade to that location; and choice trials (Figure 2B), in which two cues were presented, one on either side of the fixation spot, and the monkeys indicated their choice by making a saccade to either location. The choice trials enabled assessment of the monkeys' behavioral preference between the risky options and a safe option. Only imperative trials were used during neuronal recording, and reaction time to key release following the appearance of the green spot was measured.

The visual cues predicted the volume of juice to be delivered at the end of each trial (Figure 1D). Juice delivery was controlled by a solenoid and custom-designed software on a Macintosh IIfx computer (Apple). Juice volume was explicitly represented by the elevation of the horizontal bars. A single bar indicated that a certain volume of juice would be delivered ($p = 1$; value cues). Two bars indicated one of two possible juice volumes would be delivered ($p = 0.5$; risk cues). It was possible that the monkeys preferred cues based on the visual properties (i.e., bar height) and likewise the neuronal activity may have been sensitive to the visual properties of the cues. Thus, an additional set of risk cues (abstract fractal images) was used to control for any visual effects on behavioral and neuronal measurements, and these were pseudocounterbalanced between monkeys (Figure 1D). Also, if monkeys valued the risk cues differentially, it is possible that differential neuronal activity in risk trials might reflect a subjective value signal rather than a risk signal per se. We tested this hypothesis by comparing the activity of a given neuron between risky trials and value trials. In summary, during recording in imperative, single-cue trials, the cues were presented in two blocks: a visual control block, consisting of three risky bar cues and three risky abstract cues (Figure 1D); and a value control block, consisting of three risky bar cues and three value bar cues (Figure 1D).

Neuronal Recording

Stereotaxic coordinates (Paxinos et al., 2000) were used to locate an area on the skull that was removed to enable access to the orbitofrontal cortex. In monkey A, after recording was completed, recording sites were marked with small electrolytic lesions (20 microamperes × 20–60 s). The animal received an overdose of pentobarbital sodium (90 mg/kg, i.v.) and was perfused with 4% paraformaldehyde in 0.1 M phosphate buffer through the left ventricle of the heart. Recording positions were reconstructed from 50 μm thick, stereotaxically oriented coronal brain sections and stained with cresyl violet. As

histological reconstruction was not available for monkey B for reasons of ongoing recordings, we reconstructed recording positions approximately. In Figure 2D, for reasons of simplicity, we collapsed recording sites from both monkeys spanning roughly 5 mm in the anterior-posterior dimension onto the same coronal outline. A small number of neurons (n = 8) were recorded from the right hemisphere in animal A within the same architectonic regions as those in the left hemisphere. No significant correlations were found between neuronal positions and responses.

Data Collection

Eye movements (behavioral choice data) were monitored using an infrared eye tracking system with 5 ms resolution (ETL200; ISCAN). The percentages of risky choices were analyzed with a two-way repeated-measures ANOVA with risk and cue type (bars or abstract) as factors. Reaction times to key release were analyzed by two single linear regressions across the three value and the three risk measures. We compared the slopes (β) of the two regressions to determine the effect sizes of these two parameters.

We recorded neuronal activity during imperative trials from single orbitofrontal neurons, typically in 20–30 min sessions, using standard electrophysiological techniques including online visualization and discrimination of neuronal impulses on oscilloscopes. Timings of neuronal discharges were stored by custom-made software running on a Macintosh IIfx computer (Apple). Typically, 10–15 trials of each trial type were performed per neuron.

Data Analysis

We analyzed the activity of each neuron in specific task epochs: 0.1–0.6 s after cue onset; 0.5 s immediately preceding the saccade, key release, and reward; and 0.5 s immediately following the saccade and the key release. The saccade, key release, and reward were $\geq$1 s apart, thus warranting separate analysis epochs. Also, we assessed immediate and tardive reward responses in four nonoverlapping, consecutive postreward epochs of 0.5 s each.

The analysis of the neuronal data comprised two consecutive steps: first identifying each specific task relationship as significant change of neuronal activity related in time to a given task event, and then determining the parameter coded by these responses.

Identification of Task Relationship

First, we identified task-related responses by comparing neuronal activity in each task epoch of all risk and value trials against a 1 s control period immediately preceding the central fixation spot. We used the paired Wilcoxon test on single trials, the pair consisting of the specific task period being tested and the control period ($p < 0.05$, two-tailed). A task-related response was defined as a significant change in activity between a given task epoch and the control period.

Assessment of Coded Parameters

The second step of the analysis used linear regression models on the task-related neuronal responses identified previously by the Wilcoxon test. In this initial risk study on orbitofrontal cortex, we aimed to separate as much as possible the risk and the value information for the animal, and presented in separate trials the risk cues with the two possible reward magnitudes (and constant mean magnitude) and the value cues with only single reward magnitudes possible (and constant no-risk) (Figures 1A, 1B, and 1D). Thus, adequate regressions of neuronal responses on risk and value in separate trials required separate single linear regressions. The regressors were the variance in risk trials (9, 36, 144 $\times$ 10^{-4}) and juice volume in value trials (0.18, 0.30, 0.42 ml). The variance of a probability distribution is defined as:

$$\mathrm{var} = \sum_i p_i * (x_i - EV)^2, \text{for } i = 1 \text{ to } n$$

which, when applied to our probability distributions, is:

$$\mathrm{var} = 0.5 * (x_{low} - EV)^2 + 0.5 * (x_{high} - EV)^2$$

where p is probability of reward, x is magnitude of reward, EV is expected value [$EV = \Sigma_i (p_i * x_i)$, i is index, and n is number of elements in the distribution (outcomes)]. In a variation of this model, we used standard deviation (0.03, 0.06, 0.12 ml) as risk regressor (square root of variance) instead of variance. In a further variation, we assessed the risk responses to the bar and abstract cues in separate regressions, in order to distinguish response relationships to the risk information conveyed by the cues from their visual sensory properties.

The slopes of the linear regressions reflected the change in impulse activity per unit of risk and value (impulses/s/ml), which is a direct physiological measure of risk and value sensitivity. In an attempt to compare the effect sizes of risk and value coding, we carried out a two-tailed t test on the unsigned β values for value versus risk, using variance and standard deviation as risk regressors separately.

As a formal test for independent coding of risk and value in the population of responses, we counted the total number of significant and nonsignificant responses for risk and value and entered these counts into a 2 $\times$ 2 contingency table. We then carried out a chi-square test for independence on the contingency table.

A second regression model tested further the postreward responses in risk trials with a multiple linear regression on risk and reward magnitude. At the time of reward, in risk trials, only a single reward magnitude was delivered, and all risk trials contained information about both the risk associated with that trial and the actually delivered reward magnitude. Thus the multiple linear regression tested in the risk trials whether the same neuronal responses following reward (y values) varied with risk or reward magnitude (two x values). Reward amounts in these trials were 0.18, 0.24, 0.30, 0.33, 0.36, and 0.42 ml. This analysis enabled us to assess more directly whether the responses coded the past risk related with reward magnitude or signaled the magnitude of the reward delivered.

The results of this analysis revealed that linear risk-related responses might alternatively be interpreted as nonmonotonic reward magnitude responses. If this were the case, these neurons should also respond to reward magnitude nonmonotonically in nonrisky trials. To assess this possibility, we tested a subpopulation of these neurons with the nonrisky value cues and fit the neuronal responses in value trials with linear and quadratic functions using curve estimation regression in PASW Statistics (v18.0; IBM, Chicago, IL).

ACKNOWLEDGMENTS

We thank Peter Bossaerts for discussions, Mercedes Arroyo for technical support, and the Wellcome Trust, Behavioural and Clinical Neuroscience Institute (BCNI) Cambridge and the Human Frontier Science Program for financial support.

Accepted: September 15, 2010
Published: November 17, 2010

REFERENCES

Bach, D.R., Seymour, B., and Dolan, R.J. (2009). Neural activity associated with the passive prediction of ambiguity and risk for aversive events. J. Neurosci. *29*, 1648–1656.

Bechara, A., Damasio, A.R., Damasio, H., and Anderson, S.W. (1994). Insensitivity to future consequences following damage to human prefrontal cortex. Cognition *50*, 7–15.

Bernoulli, D. (1738). Specimen theoriae novae de mensura sortis. Comentarii Academiae Scientiarum Imperialis Petropolitanae (Papers Imp. Acad. Sci. St. Petersburg) 5, 175–192. [Translated as: Exposition of a new theory on the measurement of risk (1954). Econometrica *22*, 23–36].

Christopoulos, G.I., Tobler, P.N., Bossaerts, P., Dolan, R.J., and Schultz, W. (2009). Neural correlates of value, risk, and risk aversion contributing to decision making under risk. J. Neurosci. *29*, 12574–12583.

Clark, L., Bechara, A., Damasio, H., Aitken, M.R., Sahakian, B.J., and Robbins, T.W. (2008). Differential effects of insular and ventromedial prefrontal cortex lesions on risky decision-making. Brain *131*, 1311–1322.

Critchley, H.D., Mathias, C.J., and Dolan, R.J. (2001). Neural activity in the human brain relating to uncertainty and arousal during anticipation. Neuron *29*, 537–545.

Fiorillo, C.D., Tobler, P.N., and Schultz, W. (2003). Discrete coding of reward probability and uncertainty by dopamine neurons. Science *299*, 1898–1902.

Hayden, B.Y., and Platt, M.L. (2007). Temporal discounting predicts risk sensitivity in rhesus macaques. Curr. Biol. *17*, 49–53.

Hayden, B.Y., Nair, A.C., McCoy, A.N., and Platt, M.L. (2008). Posterior cingulate cortex mediates outcome-contingent allocation of behavior. Neuron *60*, 19–25.

Hikosaka, K., and Watanabe, M. (2000). Delay activity of orbital and lateral prefrontal neurons of the monkey varying with different rewards. Cereb. Cortex *10*, 263–271.

Hsu, M., Bhatt, M., Adolphs, R., Tranel, D., and Camerer, C.F. (2005). Neural systems responding to degrees of uncertainty in human decision-making. Science *310*, 1680–1683.

Huettel, S.A., Stowe, C.J., Gordon, E.M., Warner, B.T., and Platt, M.L. (2006). Neural signatures of economic preferences for risk and ambiguity. Neuron *49*, 765–775.

Kennerley, S.W., and Wallis, J.D. (2009). Evaluating choices by single neurons in the frontal lobe: Outcome value encoded across multiple decision variables. Eur. J. Neurosci. *29*, 2061–2073.

Kennerley, S.W., Dahmubed, A.F., Lara, A.H., and Wallis, J.D. (2009). Neurons in the frontal lobe encode the value of multiple decision variables. J. Cogn. Neurosci. *21*, 1162–1178.

Kepecs, A., Uchida, N., Zariwala, H.A., and Mainen, Z.F. (2008). Neural correlates, computation and behavioural impact of decision confidence. Nature *455*, 227–231.

Knoch, D., Gianotti, L.R., Pascual-Leone, A., Treyer, V., Regard, M., Hohmann, M., and Brugger, P. (2006). Disruption of right prefrontal cortex by low-frequency repetitive transcranial magnetic stimulation induces risk-taking behavior. J. Neurosci. *26*, 6469–6472.

Kobayashi, S., Pinto de Carvalho, O., and Schultz, W. (2010). Adaptation of reward sensitivity in orbitofrontal neurons. J. Neurosci. *30*, 534–544.

Kuhnen, C.M., and Knutson, B. (2005). The neural basis of financial risk taking. Neuron *47*, 763–770.

Levy, H., and Markowitz, H.M. (1979). Approximating expected utility by a function of mean and variance. Am. Econ. Rev. *69*, 308–317.

Levy, I., Snell, J., Nelson, A.J., Rustichini, A., and Glimcher, P.W. (2010). Neural representation of subjective value under risk and ambiguity. J. Neurophysiol. *103*, 1036–1047.

McCoy, A.N., and Platt, M.L. (2005). Risk-sensitive neurons in macaque posterior cingulate cortex. Nat. Neurosci. *8*, 1220–1227.

Mobini, S., Body, S., Ho, M.Y., Bradshaw, C.M., Szabadi, E., Deakin, J.F., and Anderson, I.M. (2002). Effects of lesions of the orbitofrontal cortex on sensitivity to delayed and probabilistic reinforcement. Psychopharmacology (Berl.) *160*, 290–298.

Padoa-Schioppa, C., and Assad, J.A. (2006). Neurons in the orbitofrontal cortex encode economic value. Nature *441*, 223–226.

Pais-Vieira, M., Lima, D., and Galhardo, V. (2007). Orbitofrontal cortex lesions disrupt risk assessment in a novel serial decision-making task for rats. Neuroscience *145*, 225–231.

Paxinos, G., Huang, X.-F., and Toga, A. (2000). The Rhesus Monkey Brain in Stereotaxic Coordinates (San Diego, CA: Academic Press).

Peters, J., and Büchel, C. (2009). Overlapping and distinct neural systems code for subjective value during intertemporal and risky decision making. J. Neurosci. *29*, 15727–15734.

Pratt, J.W. (1964). Risk Aversion in the Small and in the Large. Econometrica *32*, 122–136.

Preuschoff, K., Bossaerts, P., and Quartz, S.R. (2006). Neural differentiation of expected reward and risk in human subcortical structures. Neuron *51*, 381–390.

Preuschoff, K., Quartz, S.R., and Bossaerts, P. (2008). Human insula activation reflects risk prediction errors as well as risk. J. Neurosci. *28*, 2745–2752.

Quilodran, R., Rothé, M., and Procyk, E. (2008). Behavioral shifts and action valuation in the anterior cingulate cortex. Neuron *57*, 314–325.

Roesch, M.R., and Olson, C.R. (2004). Neuronal activity related to reward value and motivation in primate frontal cortex. Science *304*, 307–310.

Roesch, M.R., Taylor, A.R., and Schoenbaum, G. (2006). Encoding of time-discounted rewards in orbitofrontal cortex is independent of value representation. Neuron *51*, 509–520.

Rogers, R.D., Everitt, B.J., Baldacchino, A., Blackshaw, A.J., Swainson, R., Wynne, K., Baker, N.B., Hunter, J., Carthy, T., Booker, E., et al. (1999). Dissociable deficits in the decision-making cognition of chronic amphetamine abusers, opiate abusers, patients with focal damage to prefrontal cortex, and tryptophan-depleted normal volunteers: Evidence for monoaminergic mechanisms. Neuropsychopharmacology *20*, 322–339.

Roitman, J.D., and Roitman, M.F. (2010). Risk-preference differentiates orbitofrontal cortex responses to freely chosen reward outcomes. Eur. J. Neurosci. *31*, 1492–1500.

Rothschild, M., and Stiglitz, J.E. (1970). Increasing risk: I. A definition. J. Econ. Theory *2*, 225.

Sanfey, A.G., Hastie, R., Colvin, M.K., and Grafman, J. (2003). Phineas gauged: Decision-making and the human prefrontal cortex. Neuropsychologia *41*, 1218–1229.

Thorpe, S.J., Rolls, E.T., and Maddison, S. (1983). The orbitofrontal cortex: Neuronal activity in the behaving monkey. Exp. Brain Res. *49*, 93–115.

Tobler, P.N., O'Doherty, J.P., Dolan, R.J., and Schultz, W. (2007). Reward value coding distinct from risk attitude-related uncertainty coding in human reward systems. J. Neurophysiol. *97*, 1621–1632.

Tobler, P.N., Christopoulos, G.I., O'Doherty, J.P., Dolan, R.J., and Schultz, W. (2009). Risk-dependent reward value signal in human prefrontal cortex. Proc. Natl. Acad. Sci. USA *106*, 7185–7190.

Tremblay, L., and Schultz, W. (1999). Relative reward preference in primate orbitofrontal cortex. Nature *398*, 704–708.

Tversky, A., and Kahneman, D. (1992). Advances in Prospect-Theory - Cumulative Representation of Uncertainty. J. Risk Uncertain. *5*, 297–323.

von Neumann, J., and Morgenstern, O. (1944). Theory of Games and Economic Behavior (Princeton, NJ: University Press).

Wallis, J.D., and Miller, E.K. (2003). Neuronal activity in primate dorsolateral and orbital prefrontal cortex during performance of a reward preference task. Eur. J. Neurosci. *18*, 2069–2081.

Weber, E.U., Shafir, S., and Blais, A.R. (2004). Predicting risk sensitivity in humans and lower animals: Risk as variance or coefficient of variation. Psychol. Rev. *111*, 430–445.

Xue, G., Lu, Z., Levin, I.P., Weller, J.A., Li, X., and Bechara, A. (2009). Functional dissociations of risk and reward processing in the medial prefrontal cortex. Cereb. Cortex *19*, 1019–1027.