

Risk-Responsive Orbitofrontal Neurons Track Acquired Salience

Masaaki Ogawa,^{1,4,*} Matthijs A.A. van der Meer,² Guillem R. Esber,³ Domenic H. Cerri,^{1,5} Thomas A. Stalnaker,³ and Geoffrey Schoenbaum^{1,3,*}

¹Department of Anatomy and Neurobiology, University of Maryland School of Medicine, 20 Penn Street, HSF-2 S251, Baltimore, MD 21201, USA

²Department of Biology and Center for Theoretical Neuroscience, University of Waterloo, 200 University Avenue West, Waterloo, ON N2L 3G1, Canada

³NIDA-Intramural Research Program, 250 Bayview Boulevard, Baltimore, MD 21224, USA

⁴Present address: The MIT Media Laboratory, Synthetic Neurobiology Group, Massachusetts Institute of Technology, Cambridge, MA 02139, USA

⁵Present address: Department of Psychology, University of North Carolina, 118 Davie Hall Campus Box 3270, Chapel Hill, NC 27599, USA

*Correspondence: seri0722@gmail.com (M.O.), geoffrey.schoenbaum@nih.gov (G.S.)

<http://dx.doi.org/10.1016/j.neuron.2012.11.006>

SUMMARY

Decision making is impacted by uncertainty and risk (i.e., variance). Activity in the orbitofrontal cortex, an area implicated in decision making, covaries with these quantities. However, this activity could reflect the heightened salience of situations in which multiple outcomes—reward and reward omission—are expected. To resolve these accounts, rats were trained to respond to cues predicting 100%, 67%, 33%, or 0% reward. Consistent with prior reports, some orbitofrontal neurons fired differently in anticipation of uncertain (33% and 67%) versus certain (100% and 0%) reward. However, over 90% of these neurons also fired differently prior to 100% versus 0% reward (or baseline) or prior to 33% versus 67% reward. These responses are inconsistent with risk but fit well with the representation of acquired salience linked to the sum of cue-outcome and cue-no-outcome associative strengths. These results expand our understanding of how the orbitofrontal cortex might regulate learning and behavior.

INTRODUCTION

Decision making is impacted by uncertainty and risk (Bach and Dolan, 2012; D'Acromont and Bossaerts, 2008; Kahneman and Tversky, 1984; Rushworth and Behrens, 2008). Recent reports have claimed that activity in the orbitofrontal cortex (OFC) represents these variables (Kepecs et al., 2008; O'Neill and Schultz, 2010). Yet while the activity of some neurons in these studies did covary with risk or uncertainty, the overall pattern of the firing of these neurons was also consistent with the representation of acquired salience, which is elevated in response to uncertain predictors of reward (Pearce and Hall, 1980; Pearce et al., 1982).

Often risk (or the closely related concept of uncertainty) and acquired salience are confounded. However, their signatures

diverge if one compares firing to certain reward and certain nonreward. This divergence occurs because the risk associated with certain reward and nonreward, defined as the mathematical variance in the probability (or amount) of reward (D'Acromont and Bossaerts, 2008), is equivalent and in fact zero (i.e., no risk), whereas the acquired salience of certain reward is clearly higher than that of certain nonreward (Mackintosh, 1975). Thus, neurons that represent risk should fire similarly and near baseline in both conditions, while neurons that represent salience should show higher activity in anticipation of certain reward than certain nonreward. Here we tested these predictions.

RESULTS

Risk-Responsive Orbitofrontal Neurons Fail to Conform to Specific Predictions for the Representation of Risk

We trained rats in a simple odor-cued response task (Figure 1A). Each trial was initiated by illumination of a house light, after which nose poking at a central odor port resulted in presentation of a 500 ms odor cue. After termination of the odor cue, the rats were required to unpoke from the odor port and respond at a nearby fluid well to receive a sucrose reward. The reward was delivered after a 1 s delay, termed the outcome anticipation period. One of four odors was presented on each trial, associated with 100%, 67%, 33%, or 0% probability of reward. Importantly, rats had to respond at the fluid well on every trial for the task to proceed. Odor-reward associations were kept the same throughout the experiment. The rats were trained ~4 weeks prior to recording.

During recording, rats exhibited differences in movement latency and licking during the outcome anticipation period (Figures 1B and 1C), consistent with an understanding of the reward associations. Specifically, the rats moved to the fluid well after odor offset much more quickly in anticipation of possible reward (100%, 67%, and 33%) than in anticipation of certain nonreward (0%) ($p < 0.001$, Mann-Whitney U test, Figure 1B). They also licked more frequently while waiting in the fluid well on these trial types ($p < 0.001$, Mann-Whitney

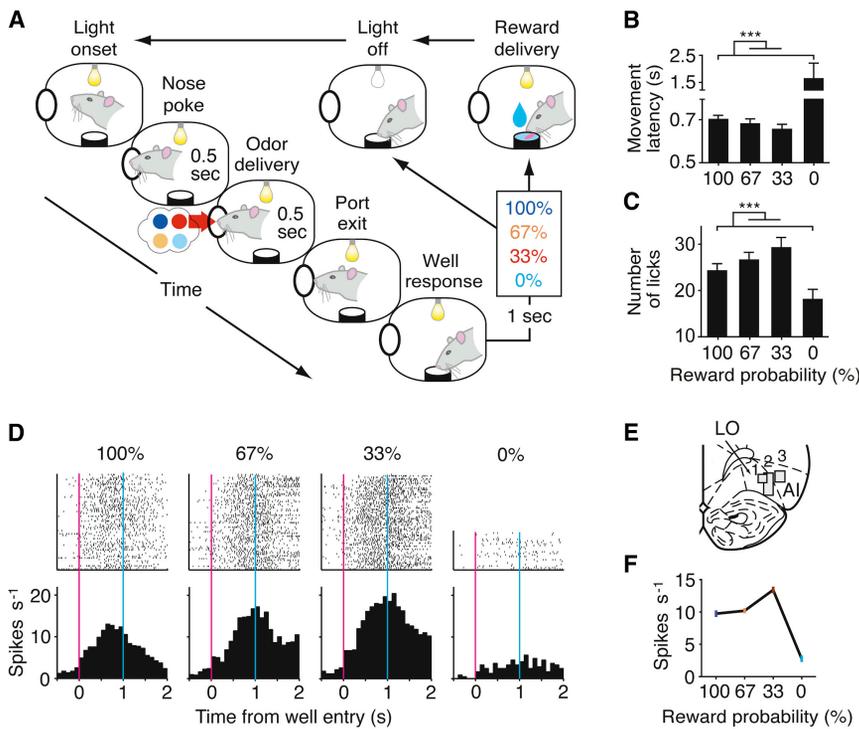


Figure 1. Task Design, Behavior Performance, and Recording from the Orbitofrontal Cortex

(A) Schematic illustrating sequence of events in the task. (B) Average latency from odor offset to responding at the fluid well ($***p < 0.001$, Mann-Whitney U test). (C) Average number of licks during 1 s after well entry ($***p < 0.001$, Mann-Whitney U test). Error bars represent SEM. (D) Raster plots and time histograms showing activity in an OFC neuron aligned to the time of response at the fluid well on trials involving 100%, 67%, 33%, or 0% probability of reward. (E) Recording sites in the lateral orbital (LO) cortex and agranular insular (AI) cortex. Shaded boxes indicate approximate extent of recording sites. We recorded 32, 222, and 78 neurons from rats 1, 2, and 3, respectively. (F) Mean firing rate of the unit shown in (D) as a function of reward probability. Rates were calculated during the outcome anticipation period (1 s) across the session. Error bars represent SEM.

U test, Figure 1C). Rats also responded more quickly and licked more as reward became increasingly uncertain (33% and 67% versus 100% and 0%, $p < 0.0001$ in movement latency; $p < 0.001$ in number of licks, Mann-Whitney U test, Figures 1B and 1C), consistent with the higher salience of these trial types. These behaviors were stable during recording (see Table S1 available online).

We recorded the activity of 332 single units in the OFC in 37 sessions in three rats (Figure 1E). Recordings were located in the lateral orbital cortex in rats 1 and 2 (total 254 neurons) and from the agranular insular cortex in rat 3 (total 78 neurons) (Figure 1E). In accord with prior reports (Kepecs et al., 2008; O'Neill and Schultz, 2010), we found a population of neurons ($n = 120/332$, 36.1%) in which firing seemed to reflect risk during the outcome anticipation period ($p < 0.05$, Mann-Whitney U test). Of these, 53 fired more while the rats were waiting for uncertain (or risky) (i.e., 33% and 67%) than certain (or nonrisky) reward (i.e., 100% and 0%) (Figures 1D, 1F, 2A, and 2C), and 67 showed the opposite pattern (Figures 2B and 2D).

Yet, while the firing of these neurons met this simple prediction for the representation of risk, their overall activity did not conform to more specific predictions of this hypothesis. For example, while these neurons fired more (or less) prior to uncertain than certain reward, they also fired differently in anticipation of certain reward (i.e., 100%) and certain nonreward (i.e., 0%; Figures 2A–2D), two conditions in which risk should have been negligible. This was reflected in the distribution of index scores comparing the firing of each neuron during the outcome anticipation period on these two trial types (i.e., 100% versus 0%), which were shifted significantly off zero ($p < 0.001$, Wilcoxon signed-rank test, Figures 2E and 2F). Indeed 76 of the 120 risk-

responsive neurons (63.3%) exhibited firing that differed significantly in anticipation of 100% and 0% reward (29 neurons for 100% > 0%, 47 neurons for 0% > 100%, $p < 0.05$, Mann-Whitney U test, Figures 2E and 2F).

Importantly, the difference in the firing to certain reward and certain nonreward typically reflected a change from baseline (1 s immediately preceding house light on) in anticipation of certain reward (34 neurons for 100% > baseline, 62 neurons baseline > 100%, $p < 0.05$); these neurons typically did not alter firing from baseline in anticipation of certain nonreward (13 neurons for 0% > baseline, 28 neurons baseline > 0%, $p < 0.05$) (Figure 2). Thus, these neurons' failure to conform to this prediction was not due to some artifact introduced by the inclusion of the 0% trials. Rather it was because these neurons showed strong responses to certain reward when risk should be negligible.

In addition, these same neurons also tended to increase (Figures 1D, 1F, 2A, 2C, and 2G) or decrease (Figures 2B, 2D, and 2H) firing during the outcome anticipation period as reward became increasingly unlikely (i.e., 33% versus 67%). This tendency did not reach significance for the neurons that suppressed firing (Figures 2D and 2H), perhaps due to the already low average firing rates (2.34 spikes/s). However, the difference was highly significant for the neurons that increased firing. The distribution of index scores comparing firing during this period on these two trial types (i.e., 33% versus 67%) was shifted significantly above zero ($p < 0.01$, Figure 2G), and nine of 53 individual neurons (17.0%) exhibited significantly higher firing prior to 33% than 67% reward.

Overall, the pattern of activity described above is inconsistent with the representation of risk, since risk should be unchanged from baseline in anticipation of certain reward and certain nonreward and equally high in anticipation of 33% and 67% reward. In fact, across the above analyses (i.e., 100% versus 0%, 33% versus 67%, and/or 100% versus baseline), only 11 of the 120

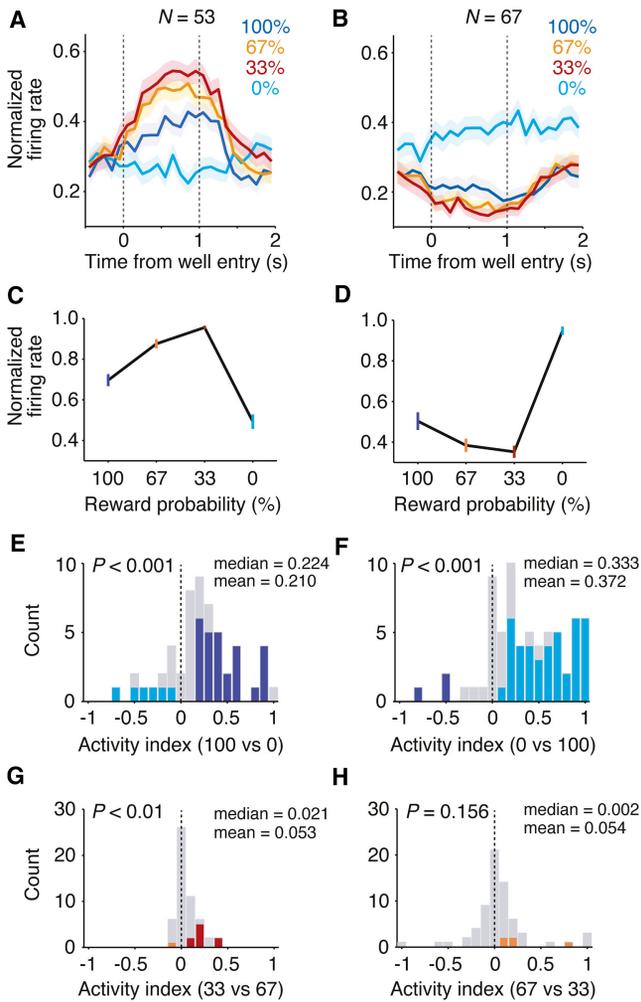


Figure 2. Risk-Sensitive Orbitofrontal Neurons Fail to Conform to Specific Predictions for Representation of Risk

(A and B) Time course of average peak-normalized firing rates in risk-responsive neurons ([33% and 67%] > [100% and 0%] or the opposite, respectively, $p < 0.05$, Mann-Whitney U test) on trials associated with different probabilities of reward (blue, 100%; orange, 67%; red, 33%; cyan, 0%) aligned to responding at the fluid well. Shading represents SEM. (C and D) Average firing during the outcome anticipation period (1 s) as a function of reward probability for 53 neurons shown in (A) and 67 neurons shown in (B), respectively. Error bars represent SEM. (E and F) Distribution of activity indices contrasting average firing in anticipation of 100% versus 0% reward during the outcome anticipation period. Activity indices were calculated as follows: (firing rate in anticipation of 100% reward [“100”] – firing rate in anticipation of 0% reward [“0”]) / (“100” + “0”) for (E) or (“0” – “100”) / (“0” + “100”) for (F) (blue bar, neurons that fired significantly more in anticipation of 100% than 0% reward; cyan bar, neurons that showed the opposite pattern, $p < 0.05$, Mann-Whitney U test). The distributions were shifted significantly above zero ($p < 0.001$, Wilcoxon signed-rank test). (G and H) Distribution of activity indices contrasting average firing in anticipation of 33% versus 67% reward during the outcome anticipation period. Activity indices were calculated as follows: (firing rate in anticipation of 33% reward [“33”] – firing rate in anticipation of 67% reward [“67”]) / (“33” + “67”) for (G) or (“67” – “33”) / (“67” + “33”) for (H) (red bar, neurons that fired significantly more in anticipation of 33% than 67% reward; orange bar, neurons that showed the opposite pattern, $p < 0.05$, Mann-Whitney U test). The distributions were shifted above zero (p value is from Wilcoxon signed-rank test).

of the apparent risk-responsive neurons met this stricter definition for the representation of risk (9.2%). Thus, while our approach identified neurons in which firing differed in anticipation of risky versus nonrisky reward, as described in earlier reports (Kepecs et al., 2008; O’Neill and Schultz, 2010), it also revealed that the firing in these neurons largely failed to conform to more specific (and quite simple) a priori predictions for the representation of risk or uncertainty in our task.

OFC Neurons Track Acquired Salience

What then might these neurons be representing? One possibility is that these neurons are tracking the heightened salience associated with both certain and uncertain reward. For decades, learning theorists have shown that the salience (i.e., the ability of a stimulus to capture attention) of reliable predictors is greater than that of poor predictors, while predictors of uncertain outcomes also acquire high levels of salience (Le Pelley, 2004; Mackintosh, 1975; Pearce and Hall, 1980; Pearce et al., 1982). Recently, Esber and Haselgrove (2011) have proposed a model that reconciles these apparently contradictory influences of predictiveness and uncertainty by proposing that an event’s salience reflects that event’s overall or combined associative strength. At the heart of this new model is the notion that the unexpected omission of reward is an emotionally potent outcome as capable of contributing to salience as reward itself (Konorski, 1967; Papini et al., 2006).

Following the Pearce-Hall model (Pearce and Hall, 1980; Pearce et al., 1982), Esber and Haselgrove (2011) assume that when a cue is probabilistically reinforced, occasional pairing of the cue and the outcome will lead to the formation of a cue-outcome or CS-US association. Once the outcome becomes expected on the basis of the cue, however, occasional omission of the outcome will encourage the formation of a second association, in this case between the cue and the emotional consequences of outcome-omission, such as frustration or disappointment. This association is referred to as the cue-no-outcome or CS-noUS association. Notably, a critical tenet of the model is that association with omission can be established only when there is association with reward, so the CS-noUS association for certain nonreward is zero.

For the purpose of determining acquired salience (ϵ), the CS-US and CS-noUS associative strengths combine to produce additive effects. At asymptote (i.e., in well-trained animals such as those in the current study), the model is reduced to:

$$\epsilon = f(\text{CS-US} + w(\text{CS-noUS}))$$

where f is a monotonically increasing function, and (CS-US + w (CS-noUS)) is combined associative strength where w is the relative weighting of these two components.

The Esber-Haselgrove model makes several interesting predictions relevant to the current data. First, a cue that predicts certain reward should have a higher salience or associative strength than a cue that predicts certain nonreward. Second, cues predictive of both reward and reward omission (i.e., cues that are reinforced probabilistically) should acquire still higher salience. And third, the acquired salience of a cue rewarded 33% of the time should be higher than that of a cue rewarded

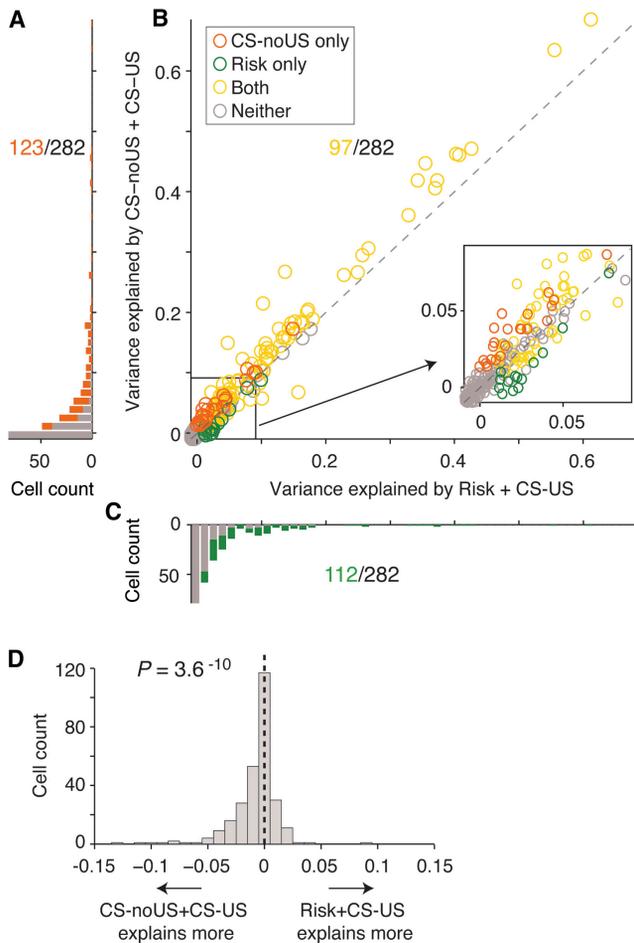


Figure 3. The Acquired Salience Model Better Explains Activity of OFC Neurons than the Risk Model

(A and C) Distribution of variance in firing rate explained (adjusted R^2) by addition of both the CS-noUS and CS-US regressors from the acquired salience model (A) or by addition of both the risk and CS-US regressors from the risk model (C), after the effects of behavior latencies and number of licks were accounted for; bar with orange (A) or green (C), count of neurons whose variance in firing was explained significantly ($p < 0.05$) by the CS-noUS (A) or the risk (C) regressor, respectively; gray bar, variance was not explained significantly; number with orange (A) or with green (C), total number of the neurons in each group. Spike counts of all task-responsive neurons were taken from the outcome anticipation period (1 s). (B) Comparison of variance in firing rate explained by the CS-noUS and CS-US regressors from the acquired salience model versus the risk and CS-US regressors from the risk model; orange or green circle, a unit whose variance in firing was explained significantly only by the CS-noUS or the risk regressor, respectively; yellow circle, a unit in which variance in firing was explained significantly both by the CS-noUS and the risk regressors; gray circle, neither; number with yellow, total number of the neurons in the corresponding group. The distribution of variance explained below 0.1 is magnified. (D) Distribution of variance explained below 0.1 is magnified. (D) Distribution of variance explained by both the CS-noUS and CS-US regressors versus that explained by both the risk and CS-US regressors. The difference was calculated for each of the task-responsive neurons as follows: (variance explained by risk and CS-US) minus (variance explained by CS-noUS and CS-US). The distribution was shifted significantly below zero ($p = 3.6 \times 10^{-10}$, Wilcoxon signed-rank test).

at 67%, at least after some experience (see [Supplemental Experimental Procedures](#) for formal description of the full model). A brief inspection of [Figure 2](#) shows that these predictions are congruent with the pattern of activity in OFC neurons.

To test more formally whether OFC neurons might track acquired salience better than risk, we fit linear regression models for acquired salience or risk to the firing of each task-responsive neuron during the outcome anticipation period (282 neurons, see [Supplemental Experimental Procedures](#) for definition). The “acquired salience” model included two critical regressors for the CS-US and CS-noUS associations, defined as $[1, 0.67, 0.33, 0]$ and $[0, 0.33, 0.67, 0]$, respectively, for [100%, 67%, 33%, 0%] probability of reward. For comparison, the “risk” model included both the CS-US regressor and a risk regressor, defined as $[0, 1, 1, 0]$. The inclusion of the CS-US regressor allowed comparison of the two models without changing the number of regressors and parameters (see [Supplemental Experimental Procedures](#) for the detail of the regression analysis).

This comparison showed that the variance explained by the acquired salience model was significantly greater than that explained by the risk model ($p = 3.6 \times 10^{-10}$, Wilcoxon signed-rank test, [Figure 3D](#)). The superiority of the acquired salience model was evident across a range of metrics, including Akaike’s information criterion (AIC), partial correlation, and leave-one-out cross-validation ([Table S2](#)). The conclusion was the same even when the analysis included all recorded neurons (332 neurons) ($p = 5.9 \times 10^{-10}$) and when log or square root transformed spike counts were used ($p < 1.0 \times 10^{-9}$).

We next computed the number of neurons in which variance in firing was significantly explained by the addition of the critical regressors that differentiated the two models: the CS-noUS regressor from the acquired salience model and the risk regressor from the risk model. This analysis identified 123 neurons in which variance in firing was explained significantly by addition of the CS-noUS regressor ([Figure 3A](#), $p < 0.05$). The activity of 97 of these neurons could also be explained by the addition of risk ([Figures 3B and 3C](#)). This overlap occurs because the CS-noUS and the risk regressors are very similar ($[0, 0.33, 0.67, 0]$ versus $[0, 1, 1, 0]$). However, consistent with our initial analysis of risk-responsive neurons, nearly all of these neurons violated predictions for the representation of risk. Indeed, only five of the 112 neurons in which variance was explained by risk in our regression analysis exhibited equivalent firing in anticipation of risky (33% versus 67%) and nonrisky (100% and 0%) reward.

The result of the regression analysis shows that when added as independent terms to a regression model, the terms in acquired salience model (CS-noUS and CS-US) outperform those in risk model (risk and CS-US) in accounting for firing in OFC neurons. However, if these neurons really represent acquired salience, rather than its component parts, then their activity should reflect the sum of the CS-US and CS-noUS associations. Consistent with this, we found activity in 129 of 282 task-responsive neurons was explained significantly by the CS-US regressor ([Figures 4A and 4B](#)), while activity in 123 neurons was explained significantly by the CS-noUS regressor ([Figures 4B and 4C](#)). These populations included 90 neurons in which activity was explained significantly by both regressors

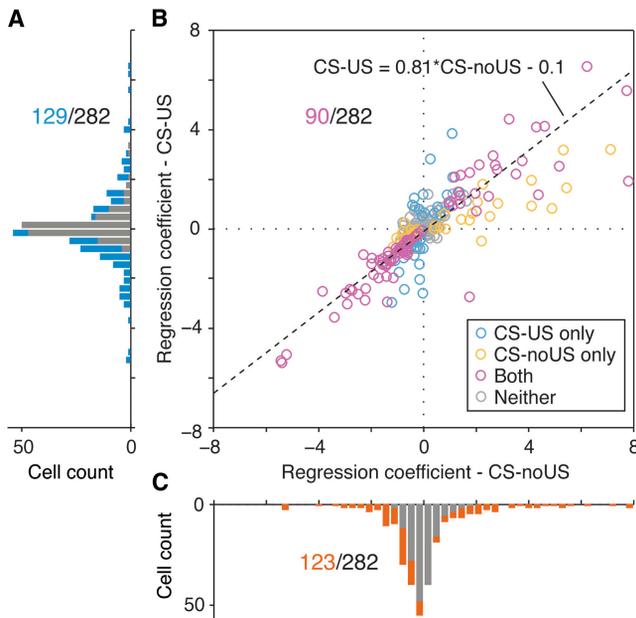


Figure 4. Acquired Salience, Reflecting the Sum of CS-US and CS-noUS Regressors, Is the Critical Factor in Explaining OFC Neurons' Firing

(A and C) Distribution of the regression coefficients of the CS-US (A) or CS-noUS (C) regressor from the acquired salience model that was fitted to the activity of each of 282 task-responsive neurons; bar with blue (A) or orange (C), count of neurons in which variance in firing were explained significantly ($p < 0.05$) by addition of the CS-US or the CS-noUS, respectively; gray bar, not significant. Total number of the neurons significantly explained by the CS-US (A) or the CS-noUS (C) regressor is shown with blue and orange, respectively. (B) Comparison of the regression coefficients of the CS-US regressor versus that of CS-noUS regressor across all task-responsive neurons (blue or orange circle, a unit whose variance in firing was explained significantly only by the CS-US or only by the CS-noUS, respectively; magenta circle, both; gray circle, neither; number with magenta, total number of the neurons in the corresponding group. Three units (11.22, 6.24), (11.83, 14.73), and (-18.28, -18.66) for (CS-noUS, CS-US), whose firing were explained significantly by both regressors, are not shown for visualization.

(chi-square test, $\chi^2 = 66.12$, $p = 4.25 \times 10^{-16}$), including 89 that had the same signs for the two regression coefficients (i.e., positive-positive or negative-negative, Figure 4B). Variance in firing explained by the addition of the CS-US and CS-noUS increased steeply in these neurons from well entry until the time of potential reward delivery, consistent with signaling of information about outcomes (Figure S1). Interestingly, the regression coefficients of the CS-noUS association were significantly larger than those for the CS-US across all 282 task-responsive neurons ($p < 0.05$, Wilcoxon signed-rank test) (i.e., $w > 1$, Figure 4B), consistent with the ordering of the specific values for acquired salience (33% > 67% > 100% > 0%) predicted by the Esber-Haselgrove model. This pattern was also apparent in the average firing rates of individual neurons (Figure S2).

DISCUSSION

That the firing of a substantial population of OFC neurons differed in anticipation of uncertain (or risky) reward versus

certain (or nonrisky) reward is consistent with prior single-unit (Kepecs et al., 2008; O'Neill and Schultz, 2010) and imaging studies (Tobler et al., 2007). However, contrasting their activity on trials in which risk or uncertainty was held constant while the likelihood of reward and nonreward varied revealed that only a very small handful of these neurons (five out of 282 task-responsive neurons, 1.8%) met a principled definition for risk encoding. While this small percentage may be compatible with that of at least one prior report (i.e., 45 responses out of 1,083 task-related responses, 4.2%) (O'Neill and Schultz, 2010), it did not rise above the level of chance in any of our analyses.

Nor was the firing of these neurons better explained by a combination of the value of the reward (i.e., CS-US) and its potential risk. While this explanation would capture the apparent reward-responsiveness of these neurons and their higher (or lower) firing to risky reward, it would also predict higher activity in anticipation of 67% than 33% reward. An inspection of Figure 2 shows that this was not the case for the risk-responsive neurons. Instead, their activity increased as reward omission became more likely. The result of the regression analysis also shows that value plus risk does not account for the data better than acquired salience.

Notably, acquired salience, as exploited here, also predicts the results obtained in previous studies that have reported single-unit correlates of risk and uncertainty in the OFC. The first is the study by O'Neill and Schultz (2010), which showed that the firing of OFC neurons covaried with the variance (or risk) in reward magnitude. Of course, shifts in the size of reward are only effective to the extent that the animal perceives them. Thus, they amount to the presentation of more or less reward than expected, on average, after a given cue. Within the context of the Esber-Haselgrove model, this is similar to what occurs when reward is presented or omitted on a probabilistic basis; in each case, the model predicts the formation of both CS-US and CS-noUS associations and, therefore, greater combined associative strength—and salience—for cues alternately reinforced with high- and low-value reward (Figure 5A). Consistent with this interpretation, in the O'Neill and Schultz study, animals showed shorter response latencies after presentation of cues associated with higher reward variance, indicating a higher level of salience. This behavior was described as risk-seeking, which is something that has been suggested of similar behaviors in other settings (McCoy and Platt, 2005; St Onge and Floresco, 2009); however, it is more easily understood not as a counterintuitive preference for risk but rather an attraction to the heightened salience of these cues.

The second report is a study by Kepecs et al. (2008), in which rats were trained on a two-choice odor mixture categorization task. On each trial, two odors (designated A and B) were presented in conjunction, mixed at one of several possible concentration ratios: 100:0, 68:32, 56:44, 44:56, 32:68, or 0:100. Trials in which the concentration of odor A was greater than odor B ($A > B$) were rewarded for visiting one of the two choice ports, whereas trials in which the reverse relation held were rewarded for visiting the other choice port. Not surprisingly, choice errors increased as a function of how similar the concentration ratios were. This pattern was reflected by firing in a subpopulation of OFC neurons, which fired the most on 56:44 trials, followed by

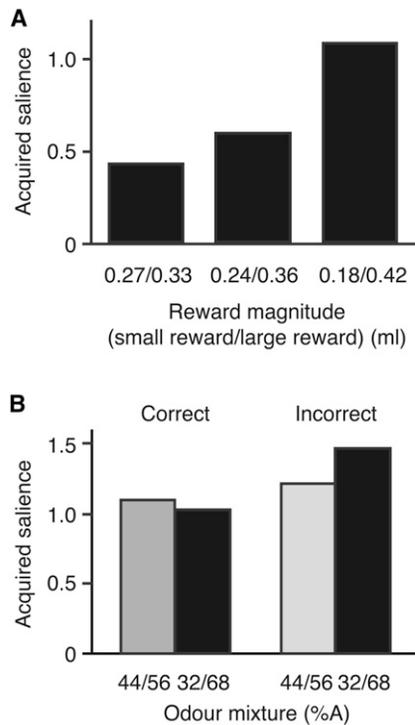


Figure 5. Acquired Salience for Different Degree of Reward Risk or Uncertainty in the Study by O'Neill and Schultz (2010) or Kepecs et al. (2008), Respectively, Simulated by the Esber-Haselgrove Model (A) Simulated acquired salience for three different cues associated with three different degrees of reward risk in O'Neill and Schultz (2010) by the full Esber-Haselgrove model. In "0.27/0.33" condition, for example, a cue was equally associated with either 0.27 ml or 0.33 ml of liquid reward. Risk for the three cues associated with 0.27/0.33, 0.24/0.36, and 0.18/0.42 are 0.0009, 0.0036, and 0.0144, respectively. (B) Simulated acquired salience for four different conditions defined by two different cues and two different behavioral responses in Kepecs et al. (2008) by the full Esber-Haselgrove model. "Correct, 44/56," for example, denotes the condition in which subjects made correct response to obtain liquid reward after presentation of a mixed odor consisting of 44% of odor A and 56% of odor B.

68:32 trials, and least on 100:0 trials. Thus, these neurons seemed to track uncertainty, either in stimulus detection or in prediction of reward. This quantity was termed decision confidence, since it was related to the rats' choice behavior, though the authors noted that the pattern was also consistent with outcome uncertainty (or risk) (Kepecs et al., 2008, Supplemental Information, p. 3).

Although decision confidence is an intriguing explanation, the Esber-Haselgrove model offers an alternative interpretation of the observed firing pattern. This interpretation rests on the assumption that variations in stimulus sampling and/or the distance between the stimulus sample and memory samples result in the formation of three rather than two different sensory precepts or states. Whenever the signal is clear, for example, in the 100:0 mixture, these variations may be assumed to have a negligible influence on stimulus encoding, and the mixture will always be correctly encoded as "A > B" (or "B > A" for 0:100). However, as the concentration ratios become more

similar, these variations will lead to increasingly frequent encoding errors, making the choice more challenging. Thus, for example, although most 68:32 trials will be correctly encoded as "A > B," the variations will cause a proportion of these trials to be encoded as "A ≈ B." On "A ≈ B" trials, the rat will be uncertain as to which port to visit and will consequently randomly distribute its responses across the two choice ports. In the context of the Esber-Haselgrove model, trials encoded as "A ≈ B" are partially reinforced and therefore should acquire a higher salience or combined associative strength than trials encoded as "A > B," which are continuously reinforced. Critically, in this behavioral setting, incorrect trials will consist entirely of highly salient "A ≈ B" trials, whereas correct trials will consist of a mixture of the highly salient "A ≈ B" trials as well as of the less salient "A > B" trials. This means that, on average, the acquired salience on incorrect trials will be higher than that on correct trials, a prediction that matches the pattern of neural activity reported by Kepecs et al. (2008) (Figure 5B).

This account is also able to explain why the difference in activity between correct and incorrect trials is less pronounced for the 56:44 than the 68:32 mixtures. Sampling variations of the more challenging 56:44 mixture will translate into a greater number of these trials being encoded as "A ≈ B." A higher proportion of "A ≈ B" trials within correct trials relative to the 68:32 mixture will make the average acquired salience on correct and incorrect trials more similar to one another. Additionally, the gap in acquired salience between correct and incorrect trials can be further bridged if it is assumed that a subset of 56:44 incorrect trials consists of (erroneously encoded) "A < B" trials. The acquired salience of "A < B" trials will be the same as that of "A > B" trials (and lower than that of "A ≈ B" trials) because the vast majority of trials encoded as "A < B" are also continuously reinforced in mixtures 0:100, 32:68, and 44:56. Therefore, the presence of erroneous "A < B" trials will lower the average salience on incorrect trials for the 56:44 mixture, bringing it even closer to the average salience on correct trials. Simulations (Figure 5B) confirmed these predictions of the Esber-Haselgrove model concerning Kepecs et al. (2008)'s results. Of course, this account cannot address all of the mechanisms that drive choice between alternatives; it simply points out that the encoding of associative relationships embedded in choice situations (as in Kepecs et al., 2008) may be sufficient to account for certain aspects of neural activity (see also Bach and Dolan, 2012).

Critically, while the data in these two reports are equally well explained by risk/uncertainty or acquired salience, this is not the case for the data in the current study, which included conditions to dissociate predictions of these two models: certain reward versus certain nonreward and 33% versus 67% reward. For both of these comparisons, risk/uncertainty predicts no difference, whereas acquired salience or combined associative strength should differ. In each case, OFC neurons that were seemingly sensitive to risky reward show a clear asymmetry in their neural activity.

Representation of acquired salience provides an account of the firing of these neurons that is also parsimonious with existing evidence that the OFC signals information about expected outcomes (Murray et al., 2007). Notably if outcome signaling in the OFC contributes to modulating attention or salience, this

would substantially expand the potential role the OFC may play in associative learning. In fact, there are already hints of such a role in existing evidence implicating the OFC in latent inhibition and formation of attentional sets, two behaviors often taken as cardinal evidence for experiential declines in salience (Chase et al., 2012; Schiller and Weiner, 2004; Schiller et al., 2006). Additionally, it has been shown recently that the OFC is required for rapid learning when cue-outcome relationships are changing by enabling the credit for a particular outcome to be assigned to a specific choice of a cue that leads to the outcome (“credit assignment”) (Walton et al., 2010). If salience were not being properly allocated, this might result in improper credit assignment in such changeable environment.

Of course, like claims that OFC neurons signal risk, uncertainty, or even decision confidence, the proposal that the OFC signals acquired salience is a hypothesis that needs to be further tested. However, this hypothesis does make concrete predictions. For example, similar signaling may be evident in downstream areas, a prediction that has recently been met for the temporoparietal junction (Kahnt et al., 2012; Kahnt and Tobler, 2013). Additionally, if the OFC contributes to experiential increases in salience, then the OFC should be necessary for the more rapid discrimination learning observed when cues have previously been paired with uncertain reward (Haselgrove et al., 2010). Similarly, given the observation of increases in firing for both higher and lower salience, OFC might also be important for behaviors that reflect a decline in salience. Indeed this may be reflected in the existing results implicating the OFC in latent inhibition and set formation described above. Our data may provide a neurophysiological substrate to explain these results. This explanation would have relevance to understanding the involvement of OFC dysfunction in neuropsychiatric disorders that involve aberrant behavior and learning, such as obsessive-compulsive disorder (Rauch et al., 1994), schizophrenia (Corlett et al., 2007), and addiction (Lucantonio et al., 2012).

EXPERIMENTAL PROCEDURES

Subjects

Male Long-Evans rats (Charles River, 350–450 g) were housed individually on a 12 hr light/dark cycle with ad libitum access to food. Water was restricted to that earned as reward in the task and to ~10 min free access after each testing session. Procedures were conducted at the University of Maryland School of Medicine in accordance with University and NIH guidelines.

Surgical and Single-Unit Recording Procedures

Procedures were as described previously (Roesch et al., 2006), except that the current experiment utilized drivable stereotrodes (Ramus and Eichenbaum, 2000). Bundles were implanted at 3.6 mm anterior to bregma, 3.2 mm laterally, and 3.2 mm ventral to the brain surface. Prior to implantation, wires were cut with surgical scissors to extend ~1.5 mm beyond the cannula and electroplated with platinum (H₂PtCl₆, Aldrich) to an impedance of ~300 kOhm. Cephalexin (15 mg/kg per oral) was administered twice daily for 2 weeks postoperatively to prevent infection. At the end of recording, the final electrode position was marked. The rats were euthanized and their brains were processed for histology using standard techniques. Neural activity was recorded using Multichannel Acquisition Processor systems (Plexon). Waveforms (>2.5:1 signal-to-noise) were extracted from active channels and recorded to disk by an associated workstation with event timestamps from the behavior computer. Units were sorted later using Offline Sorter software (Plexon). Sorted files were then processed in Neuroexplorer

to extract unit timestamps and relevant event markers for analysis in MATLAB 2012a (MathWorks).

Behavioral Task

Recording was conducted in aluminum chambers 18” on each side. A central odor port was located above a fluid well connected to an air flow dilution olfactometer to allow the rapid delivery of olfactory cues. Task control was implemented via computer; port entry and licking was monitored by photo-beams. One of four different odors (Auralva, Para-isopropyl hydratropic aldehyde, Camekol DH, or Verbena olifac) was presented on each trial, associated with 100%, 67%, 33%, or 0% probability of reward. After odor offset, rats were required to make a response at the fluid well within 100 s. In other words, rats had to wait 100 s for the task to proceed if they did not respond at the fluid well. In a rewarded trial, a 0.1 ml bolus of 10% sucrose solution was delivered 1 s after well entry. The house light was turned off when rats left the fluid well after reward consumption. In a nonrewarded trial, the house light was turned off 1 s after the entry. Rats did not need to wait for 1 s in the well; all that was required for the task to proceed was a response at the well. Rats were trained for ~4 weeks prior to the start of recording, such that they responded on all rewarded trials and >99.0% of the trials in which the odor associated with 0% reward is presented.

SUPPLEMENTAL INFORMATION

Supplemental Information includes two figures, two tables, and Supplemental Experimental Procedures and can be found with this article online at <http://dx.doi.org/10.1016/j.neuron.2012.11.006>.

ACKNOWLEDGMENTS

We thank Reza Ramezan for comments on the statistical analysis. This work was supported by funding from the Uehara Memorial Foundation to M.O., the National Science and Engineering Research Council and the Netherlands Organization for Scientific Research (VENI 863.10.013) to M.A.A.vdM., and NIDA to G.S. and M.O. This work was conducted while G.S. and T.A.S. were employed at the University of Maryland. The opinions expressed in this article are the authors’ own and do not reflect the view of the National Institutes of Health, the Department of Health and Human Services, or the United States government. M.O. and G.S. conceived the experiments; M.O. and D.H.C. carried out the experiment. M.O. and M.A.A.vdM. analyzed the data with assistance from G.R.E., T.A.S., and G.S. The manuscript was prepared by M.O. and G.S. with assistance from M.A.A.vdM. and G.R.E.

Accepted: November 3, 2012

Published: January 23, 2013

REFERENCES

- Bach, D.R., and Dolan, R.J. (2012). Knowing how much you don’t know: a neural organization of uncertainty estimates. *Nat. Rev. Neurosci.* 13, 572–586.
- Chase, E.A., Tait, D.S., and Brown, V.J. (2012). Lesions of the orbital prefrontal cortex impair the formation of attentional set in rats. *Eur. J. Neurosci.* 36, 2368–2375.
- Corlett, P.R., Murray, G.K., Honey, G.D., Aitken, M.R., Shanks, D.R., Robbins, T.W., Bullmore, E.T., Dickinson, A., and Fletcher, P.C. (2007). Disrupted prediction-error signal in psychosis: evidence for an associative account of delusions. *Brain* 130, 2387–2400.
- D’Acromont, M., and Bossaerts, P. (2008). Neurobiological studies of risk assessment: a comparison of expected utility and mean-variance approaches. *Cogn. Affect. Behav. Neurosci.* 8, 363–374.
- Esber, G.R., and Haselgrove, M. (2011). Reconciling the influence of predictiveness and uncertainty on stimulus salience: a model of attention in associative learning. *Proc. Biol. Sci.* 278, 2553–2561.

- Haselgrove, M., Esber, G.R., Pearce, J.M., and Jones, P.M. (2010). Two kinds of attention in Pavlovian conditioning: evidence for a hybrid model of learning. *J. Exp. Psychol. Anim. Behav. Process.* *36*, 456–470.
- Kahneman, D., and Tversky, A. (1984). Choices, values, and frames. *Am. Psychol.* *39*, 341–350.
- Kahnt, T., and Tobler, P. (2013). Salience signals in the right temporoparietal junction facilitate value-based decisions. *J. Neurosci.* Published online January 16, 2013. <http://dx.doi.org/10.1523/JNEUROSCI.3531-12.2013>.
- Kahnt, T., Chang, L.J., Park, S.Q., Heinzle, J., and Haynes, J.-D. (2012). Connectivity-based parcellation of the human orbitofrontal cortex. *J. Neurosci.* *32*, 6240–6250.
- 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.
- Konorski, J. (1967). *Integrative Activity of the Brain* (Chicago: University of Chicago Press).
- Le Pelley, M.E. (2004). The role of associative history in models of associative learning: a selective review and a hybrid model. *Q. J. Exp. Psychol.* *57*, 193–243.
- Lucantonio, F., Stalnaker, T.A., Shaham, Y., Niv, Y., and Schoenbaum, G. (2012). The impact of orbitofrontal dysfunction on cocaine addiction. *Nat. Neurosci.* *15*, 358–366.
- Mackintosh, N.J. (1975). A theory of attention: variations in the associability of stimuli with reinforcement. *Psychol. Rev.* *82*, 276–298.
- McCoy, A.N., and Platt, M.L. (2005). Risk-sensitive neurons in macaque posterior cingulate cortex. *Nat. Neurosci.* *8*, 1220–1227.
- Murray, E.A., O'Doherty, J.P., and Schoenbaum, G. (2007). What we know and do not know about the functions of the orbitofrontal cortex after 20 years of cross-species studies. *J. Neurosci.* *27*, 8166–8169.
- O'Neill, M., and Schultz, W. (2010). Coding of reward risk by orbitofrontal neurons is mostly distinct from coding of reward value. *Neuron* *68*, 789–800.
- Papini, M.R., Wood, M., Daniel, A.M., and Norris, J.N. (2006). Reward loss as a psychological pain. *Int. J. Psychol. Psychol. Ther.* *6*, 189–213.
- Pearce, J.M., and Hall, G. (1980). A model for Pavlovian learning: variations in the effectiveness of conditioned but not of unconditioned stimuli. *Psychol. Rev.* *87*, 532–552.
- Pearce, J.M., Kaye, H., and Hall, G. (1982). Predictive accuracy and stimulus associability: development of a model for Pavlovian learning. In *Quantitative Analyses of Behavior*, M.L. Commons, R.J. Herrnstein, and A.R. Wagner, eds. (Cambridge, MA: Ballinger), pp. 241–255.
- Ramus, S.J., and Eichenbaum, H. (2000). Neural correlates of olfactory recognition memory in the rat orbitofrontal cortex. *J. Neurosci.* *20*, 8199–8208.
- Rauch, S.L., Jenike, M.A., Alpert, N.M., Baer, L., Breiter, H.C., Savage, C.R., and Fischman, A.J. (1994). Regional cerebral blood flow measured during symptom provocation in obsessive-compulsive disorder using oxygen 15-labeled carbon dioxide and positron emission tomography. *Arch. Gen. Psychiatry* *51*, 62–70.
- 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.
- Rushworth, M.F., and Behrens, T.E. (2008). Choice, uncertainty and value in prefrontal and cingulate cortex. *Nat. Neurosci.* *11*, 389–397.
- Schiller, D., and Weiner, I. (2004). Lesions to the basolateral amygdala and the orbitofrontal cortex but not to the medial prefrontal cortex produce an abnormally persistent latent inhibition in rats. *Neuroscience* *128*, 15–25.
- Schiller, D., Zuckerman, L., and Weiner, I. (2006). Abnormally persistent latent inhibition induced by lesions to the nucleus accumbens core, basolateral amygdala and orbitofrontal cortex is reversed by clozapine but not by haloperidol. *J. Psychiatr. Res.* *40*, 167–177.
- St Onge, J.R., and Floresco, S.B. (2009). Dopaminergic modulation of risk-based decision making. *Neuropsychopharmacology* *34*, 681–697.
- 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.
- Walton, M.E., Behrens, T.E.J., Buckley, M.J., Rudebeck, P.H., and Rushworth, M.F.S. (2010). Separable learning systems in the macaque brain and the role of orbitofrontal cortex in contingent learning. *Neuron* *65*, 927–939.