

Published in final edited form as:

Neuron. 2013 March 20; 77(6): 1174–1186. doi:10.1016/j.neuron.2013.01.008.

Dopaminergic reward signals selectively decrease fMRI activity in primate visual cortex

J.T. Arsenault^{1,2}, K. Nelissen¹, B. Jarraya^{1,2}, and W. Vanduffel^{1,2,3,*}

¹Laboratory of Neuro-Psychophysiology, KU Leuven Medical School, Campus Gasthuisberg, 3000 Leuven, Belgium.

²Athinoula A. Martinos Center for Biomedical Imaging, Massachusetts General Hospital, Charlestown, MA 02129, USA.

³Department of Radiology, Harvard Medical School, Charlestown, MA 02129, USA.

Summary

Stimulus-reward coupling without attention can induce highly specific perceptual learning effects, suggesting that rewards trigger selective plasticity within visual cortex. Additionally, dopamine-releasing events - temporally-surrounding stimulus-reward associations - selectively enhance memory. These forms of plasticity may be evoked by selective modulation of stimulus representations during dopamine-inducing events. However, it remains to be shown whether dopaminergic signals can selectively modulate visual cortical activity. We measured fMRI activity in monkey visual cortex during reward-only trials apart from intermixed cue-reward trials. Rewards without visual stimulation selectively decreased fMRI activity within the cue representations that had been paired with rewards during other trials. Behavioral tests indicated that these same uncued reward trials strengthened cue-reward associations. Furthermore, such spatially-specific activity modulations depended on prediction error, as shown by manipulations of reward magnitude, cue-reward probability, cue-reward familiarity, and dopamine signaling. This cue-selective negative reward signal offers a mechanism for selectively gating sensory cortical plasticity.

Introduction

Coupling a visual stimulus with a reward improves stimulus detection (Engelmann et al., 2009; Engelmann and Pessoa, 2007), increases stimulus selection (Pessiglione et al., 2008; Pessiglione et al., 2006; Serences, 2008), and reduces reaction times (Nomoto et al., 2010; O'Doherty et al., 2004; Roesch and Olson, 2004). Furthermore, stimulus-specific perception has been enhanced by stimulus-reward coupling in the absence of attention (Seitz et al., 2009). This indicates that rewards may help regulate selective plasticity within the visual representation of reward-predicting stimuli. Nonetheless, the neural mechanisms by which rewards induce stimulus selective modulation of activity in visual cortex remain unknown.

The dopaminergic neuromodulatory system is a potential candidate for distributing reward information to visual cortex (Tan, 2009). This system is controlled by midbrain

© 2013 Elsevier Inc. All rights reserved.

*Corresponding Author: (wim@nmr.mgh.harvard.edu; +32 16 3 30666) .

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

dopaminergic neurons, which, in addition to other response properties (Fiorillo et al., 2003; Ljungberg et al., 1992; Matsumoto and Hikosaka, 2009), exhibit a phasic prediction error (PE) response signaling the difference between outcome and expectation (Bromberg-Martin et al., 2010; Schultz et al., 1997). Moreover, PE signals originating in ventral midbrain neurons are relayed through a widespread network of connections (Lidow et al., 1991; Lindvall et al., 1974), resulting in increased dopamine release (Gonon, 1988; Zhang et al., 2009), activity modulation (Pessiglione et al., 2006), and plasticity (Surmeier et al., 2010) at projection sites. Accordingly, a recent human fMRI study has shown that reward information was present throughout most brain regions tested (Vickery et al., 2011). Therefore, the highly selective behavioral and neural effects induced by stimulus-reward pairings must be reconciled with the apparent widespread and diffuse nature of neuromodulatory reward signals.

A potential explanation for this seeming contradiction is that selectivity arises through an interaction between a broadly distributed reward signal and coincident bottom-up, cue-driven activity. In this way, a diffuse dopaminergic reward signal is rendered selective, allowing rewards to specifically modulate activity within reward-predicting cue representations (Roelfsema et al., 2010; Seitz and Watanabe, 2005). In agreement with this interpretation, the pairing of an auditory stimulus with microstimulation of the ventral tegmental area (VTA), a surrogate for reward, specifically enhanced the representation of a stimulation-paired frequency within rat auditory cortex in a dopamine-dependent manner (Bao et al., 2001). In addition, Pleger et al. (2009) has found a stimulus-selective, dopaminergic reward feedback signal within somatosensory cortex.

Surprisingly though, direct evidence for selective reward modulations in primate visual cortex has not yet been demonstrated. This is probably due to the difficulty of disentangling reward from other co-occurring cognitive factors such as attention (Maunsell, 2004). For example, while Serences (2008) found that the association of a visual stimulus with a higher reward probability resulted in stimulus-selective increases in fMRI activity, the contributions of reward and attention to these results are indistinguishable. Weil et al., (2010) also looked at the effects of direct stimulus-reward relationships in visual cortex. In an effort to isolate reward effects from attention, they temporally disassociated reward from stimulus presentation. This study, however, found only a main effect of reward outside the representation of the visual stimulus suggesting these reward modulations were stimulus-specific.

In order to differentiate the contributions of attention and reward, we developed a paradigm for investigating cue-selective reward modulations that were temporally separated from discrete cue-reward association trials. The presence of these modulations is suggested by experiments in which the memory of a cue-reward association is facilitated by temporally-separated rewards or dopamine-inducing events. This form of enhancement has been demonstrated when an association is followed by sucrose consumption (Messier and White, 1984), brain stimulation rewards (White and Major, 1978), systemic amphetamine injection (Blais and Janak, 2007; Osofsky et al., 1988), amygdala injections of a D3 agonist (Hitchcott and Phillips, 1998) and exposure to novel, dopamine-inducing environments (Wang et al., 2010). Although never shown directly, the specificity of these positive behavioral effects indicates that diffuse dopaminergic reward signals preferentially modulate previously rewarded cue-representations. We therefore hypothesized that the interaction of cue and reward-driven signals not only causes selective modulation of the stimulus representation but also “tags” this representation. Subsequent dopaminergic reward modulations then interact with these “tags”, directly affecting the stimulus representation during events outside the actual cue-reward association.

To test for selective modulations in visual cortex during rewards temporally separated from stimulus-reward associations, we used a factorial paradigm with functional magnetic resonance imaging (fMRI) in monkeys (visual cue X reward) and focused on trials in which juice rewards were not cued by the visual stimulus. As hypothesized, we found spatially-specific reward modulations in the absence of visual stimulation. Manipulations of reward magnitude, cue-reward probability, and cue-reward familiarity confirmed that this signal was affected by PE while concurrently excluding the possibility that other extraretinal factors -such as attention, expectation, anticipation or trial structure (Sirotin and Das, 2009)- contributed to this novel reward signal in visual cortex. Next, a pharmacological challenge showed that the reward modulation in visual cortex was controlled at least partially by dopaminergic signaling. Lastly, we demonstrated that rewards temporally separated from stimulus-reward association events positively influence the behavioral preferences of monkeys for that stimulus.

Results

Effects of Rewards on Occipital Cortical Activity in the Absence of a Visual Stimulus (Uncued Reward Activity; Experiment 1)

Our first experiment (2-by-2 factorial design) was designed to probe for the existence of reward modulations in visual cortex in the absence of visual stimulation during trials temporally separated from cue-reward association events. Monkeys were trained to fixate on a central fixation point and to wait a random interval (3.5 – 6 s) for one of four equiprobable events to occur (Figure 1A). During half of the trials, a visual cue (a green abstract shape presented for 500 ms, see Figure S1A) signaled both the end of the wait period and a 50% probability of an impending 0.2 ml juice reward (cue-reward trial, Figure 1B). Due to the temporal uncertainty generated by the randomized wait period, the visual cue indicated an immediate increase in the probability of an upcoming reward. The uncued trials (50% rewarded) conserved the average timing between trial onset and reward (3.9 – 6.4 s) but lacked the cue marking reward availability. Therefore, uncued rewards generated a larger PE than cued rewards because the administration of these rewards was not signaled by previous events. Uncued trials in which the reward was omitted (i.e. fixation trials) were used to determine baseline activity. Significantly, the design included cue-reward trials (to maintain a cue-reward association) and uncued reward trials (to test for reward-induced modulations in visual cortex without visual stimulation).

Three monkeys performed the 2-by-2 factorial design task during fMRI acquisition. Figure 2A depicts fMRI activity during uncued reward trials ($P < 0.05$, family-wise error (FWE) corrected, uncued reward minus fixation; no visual stimuli presented during either trial type) overlaid onto a flattened representation of the left occipital cortex. Surprisingly, the modulation of fMRI activity induced by the uncued reward was largely negative. Analysis of the fMRI timecourses within the cue representation (in visual areas V3, V4 and TEO) showed that the fMRI percent signal change (PSC) between the uncued reward and fixation conditions peaked at ~4 seconds after event onset (Figure S2, see Supplemental Experimental Procedures), indicating that the deactivations were associated with reward delivery. In addition, this reward-induced decrease in the fMRI activity co-localized surprisingly well with the cue-representation as determined in an independent localizer experiment (Figure 2B and 2C). To characterize the relationship between reward- and cue-driven activity, we calculated the correlation between the beta-values of these two signals voxel-by-voxel in 6 visual regions of interest (ROIs) (e.g. for V4 in Figure 2D, Supplemental Experimental Procedures). Significant correlations between cue and reward activities were found in areas V3, V4, and TEO (Figure 2E) indicating that the voxels best activated by the cue showed the strongest deactivations during uncued rewards.

We next examined the cued reward trials, which allowed us to determine whether differences in PE between cued and uncued rewards affected the magnitude of the reward modulations. Reward modulations during cued trials found within the cue-representation were negative (Figure 3A) and largely confined to the stimulus representation and were thus qualitatively similar to the reward modulations observed during the uncued conditions. We then compared the magnitude of reward modulations during the cued trials (smaller PE) and the uncued trials (larger PE). Reward modulations were found to be significantly stronger within the cue representation during the uncued rewards trials (Figure 3B) when the prediction error was larger, suggesting that the strength of the observed reward modulations depends on PE.

Are cue events necessary to evoke a negative reward signal in occipital cortex (Experiment 2)?

In experiment 2, we tested the hypothesis that the deactivations observed in visual cortex during uncued rewards were governed by the interaction between stimulus and reward during cue-reward association trials. In other words, was the presence of cued trials necessary for the deactivations observed during uncued reward trials? To achieve this, two different monkeys (M22 and M23), who were naïve with respect to the stimuli used, performed a variant of experiment 1 that consisted solely of fixation and uncued reward trials –hence, without cued trials. Within this paradigm, uncued reward activity, as monitored by a ROI analysis within the cue-representation (measured during an independent localizer scan), showed no significant reduction in activity (Figure 4, Figure S3). These results suggest that the deactivations observed during uncued reward trials in experiment 1 require the presence of randomly intermixed cue-reward trials.

Effect of Reward Size on Uncued Reward Activity (Experiment 3)

We hypothesized that by manipulating PE during uncued rewards through changes in reward size, we could alter the strength of the reward modulations in visual cortex. Importantly, the use of different reward sizes allowed us to examine the dependence of reward modulation on PE in the absence of visual stimulation without the need to compare rewarded trials to unrewarded ones (e.g. uncued reward vs. fixation). Hence, we could also rule out the possibility that the perception of “reward omission” during unrewarded trial types (fixation and cued trials) accounted for the activity modulations observed. To test the effect of reward size on reward modulations in experiment 3, we replaced the single reward level (0.2 ml) used in experiment 1 with large (0.3 ml) and small (0.1 ml) rewards. Consistent with electrophysiological studies (Tobler et al., 2005), reward-responsive regions in the ventral midbrain, presumably corresponding to the VTA, displayed stronger responses for larger unpredicted rewards (Figures 5A and 5D). The fMRI responses within the cue-representation also showed stronger deactivations associated with larger uncued rewards (Figures 5A and 5D). These differences cannot be explained by visual stimulation, as no visual cues were presented during either trial type. Furthermore, a reward omission signal cannot account for this effect as both trial types were rewarded.

In addition, we observed substantial colocalization between voxels more strongly deactivated by larger uncued rewards and voxels representing the cue (Figures 5B and 5E). We quantified the dependency of the effect of reward size (large vs. small uncued rewards) upon cue localizer activity by calculating the voxel-by-voxel correlation between the beta values of these two signals. We found a significant correlation between the two (Figure S4), confirming that the strongest deactivations evoked by administering the larger uncued reward were most prevalent within those voxels best driven by the visual cue. Interestingly, we also observed a run-by-run correlation between the activity within the cue-representation and that in the ventral midbrain during the large and small uncued reward trials (Figures 5C

and 5F, see Supplemental Experimental Procedures), which suggests that the ventral midbrain may initiate the deactivations observed in visual cortex.

Effect of Cue-Reward Probability on Uncued Reward Activity (Experiment 4)

Experiments 1 and 2 suggest that the strength of deactivations during uncued rewards depend on attributes of the cue-reward association, as does PE. Therefore, we hypothesized that representations of cues associated with higher reward probabilities would show stronger deactivations during uncued rewards, due to the increased PE response exhibited by dopaminergic neurons when a cue is associated with a higher probability of reward (Fiorillo et al., 2003). We tested this prediction in experiment 4 by manipulating the probability of reward associated with visual cues. This design used two separate cues (see Figure S1A and S1B) to examine the specificity of the uncued reward activity for the two distinct cue-representations. Initially, one cue was assigned a high reward-probability (66% of trials rewarded) and a second cue, a low reward-probability (33% of trials rewarded) (green high reward-probability example; Figure 6A). After training and scanning with this cue-reward contingency, the relationship was reversed and a second scan period began (Figure 6C and 6D). Note that although we manipulated the probability of reward associated with the visual cues, we monitored fMRI activity during uncued rewards. As hypothesized, deactivations during uncued rewards within the representation of the green cue were significantly stronger when the green cue held a high reward-probability, and vice-versa for the red cue (Figure 6B). Thus, uncued reward activity in visual cortex is sensitive to the probability of reward associated with a given cue, thereby simultaneously and differentially modulating fMRI activity within two cue-representations.

Examination of the maps of uncued reward activity generated during the green and red high reward-probability experiments show stronger deactivations within the representation of the more frequently rewarded cue (Figure S5A). In addition, one can also see a substantial overlap in the deactivation patterns generated during the two experiments. This is to be expected as there are many voxels driven by both stimuli and therefore stimulus-driven activity in these voxels co-occurs with reward delivery in both green and red high-value experiments. Despite this overlap, we asked whether the overall pattern of uncued reward activity within higher visual regions (V3-TEO) was similar to that induced by the high reward-probability stimulus. To determine this, we trained a multivariate pattern analysis (MVPA) classifier, using data from the independent localizer experiment, to distinguish between red and green cue presentations. The uncued reward activity maps were then inverted for comparison with cue localizer activity and the classifier was tested on this uncued reward activity (i.e. in the absence of visual stimulation). The classifier successfully identified the high reward-probability cue during both the green and red high reward-probability experiments (Figure S5B, see Supplemental Experimental Procedures). Thus, the pattern of activity generated by the uncued reward held information surprisingly similar, albeit of opposite polarity, to that of the visual response to the high-value stimulus itself.

The PE response of ventral midbrain dopaminergic neurons to a cued reward is stronger during the acquisition of novel contingencies (Hollerman and Schultz, 1998). Therefore, if the PE response during the cued rewards influences uncued reward activity, one would predict larger deactivations during uncued rewards directly after a reversal of cue-reward contingencies, because the relationships being learned are novel. In an effort to determine how the strength of the reward modulation changed as a function of time within experiment 4, we divided the uncued reward activity into early, middle and late time-bins for both the first and second scan periods. A cue selectivity index was then calculated, comparing reward activity within the two cue-representations at each time point (see Supplemental Experimental Procedures). The selectivity index exhibited a preference for the high-reward cue within all time-bins during the first scan period (Figure 6E and 6F), confirming the

analysis shown in Figure 6B. In addition, both animals displayed the highest selectivity during the earliest time-bin of the second scan period, immediately after the change in the cue-reward relationships (between time-bins c and d). Thus, exactly as predicted, the uncued reward modulation is strongest directly after the reversal in reward-probability, when novel contingencies are being learned. The selectivity diminished over the next two phases of the experiment (time-bins e and f), as the new cue-reward contingencies became more familiar, resulting in a significant difference in selectivity between the time-bin immediately after switching the reward probabilities and the subsequent time-bins. These results indicate that the amount of deactivation during uncued rewards is also contingent upon the level of PE during the cued reward and is therefore sensitive to familiarity with cue-reward relationships.

Effect of Cue-Reward Familiarity on Uncued Reward Activity (Experiment 5)

To corroborate these results, experiment 5 directly tested the dependence of deactivations during uncued rewards upon familiarity with cue-reward relationships (Hollerman and Schultz, 1998). We therefore used absolute cue-reward relationships (with one cue always rewarded while the second one was never rewarded; the rewarded cues were counterbalanced across animals) to examine whether exposure to these consistent associations reduced the magnitude of deactivations during uncued rewards. As hypothesized, time-bins of uncued-reward fMRI activity within the representation of the high-reward cue exhibited significant familiarity effects for the predictable cue-reward contingency, with the weakest modulations occurring within the last time-bin for either animal (Figure 7). Closer examination of the timecourse of the uncued reward activity revealed two distinct phases (Figure S6). The early phase was marked by a trend towards stronger deactivations, while the later phase displayed a significant decrease in deactivation strength as a function of cue-reward exposure. These findings show that after an initial period, the deactivations elicited by uncued rewards become reduced in strength as subjects are increasingly exposed to absolute cue-reward contingencies.

Dopaminergic Modulation of Uncued Reward Activity (Experiment 6)

Based on the earlier results, we hypothesized that the influence of the PE on visual cortical activity during uncued rewards depends upon dopaminergic signaling. To test this premise, experiment 6 examined the effects of a dopamine (D1) antagonist SCH-23390 challenge on neural activity during uncued rewards (see Figure S7 for details). Initial scans without SCH-23390 were used to monitor the baseline fMRI activity during uncued rewards (i.e. baseline phase). Afterwards, an “injection run” was performed in which two equivalent boluses (0.0025 - 0.0050 mg/kg) of SCH-23390 were administered intravenously 5 minutes apart. The effect of SCH-23390 was then monitored during the post-injection phase, followed by the recovery phase. The normalized visual response (cue – fixation) was used to test for aspecific drug effects on the fMRI response, and with the small doses utilized here, no significant effect was found across drug phases (group-level analysis, 30 runs/phase, M19 & M20 - 15 runs/phase, Kruskal-Wallis non-parametric ANOVA across phases; $p = 0.66$). In contrast, within this same group of runs, a significant drug effect was found on normalized uncued reward activity within the cue-representation (Figure 8A, see Supplemental Experimental Procedures). The diminished uncued reward signal in visual cortex measured during the post-injection phase, relative to baseline and recovery (Figure 8B-8D), shows that the amplitude of deactivations during uncued rewards depends upon dopamine signaling.

The Effect of Uncued rewards on the Strength of Cue-reward Associations (Experiment 7)

Interleaved uncued rewards may weaken cue-reward associations since rewards are not fully contingent with the cue. This leads to the hypothesis that uncued reward modulations may

represent an “unlearning” signal. Alternatively, uncued rewards may strengthen cue-reward associations. This would be in agreement with previous studies demonstrating a role for temporally-separated dopamine inducing events in strengthening cue-reward associations. Therefore, we tested behaviorally whether the strength of cue-reward associations, as measured by changes in stimulus preference, was affected by intermixed uncued reward trials. A free-choice saccade task was used to determine stimulus preference. The animals fixated centrally to begin a trial, and after a delay period (1000 - 1500 ms) two peripheral stimuli were displayed simultaneously. The monkeys had to saccade to one of the stimuli to complete a trial. Importantly, the stimulus position and the probability of being rewarded were equalized between the two stimuli, and therefore differences in stimulus selection were interpreted as a bias for a stimulus, or stimulus preference. After a baseline preference test, the animals were exposed to cue-reward association blocks, containing 25 cue-reward association trials, during which a juice reward was paired with the initially non-preferred stimulus. There were two variants of the cue-reward association blocks, those that contained uncued reward trials and those that did not (see Experimental Procedures). After the cue-reward association block, the monkey’s stimulus preference was tested again. We found a larger increase in the preference for the reward-associated cue after association sessions that contained the uncued reward trials compared to those that did not (Figure 9A and 9B). These results demonstrate that, in addition to modulating fMRI activity within the cue representation, uncued rewards temporally surrounding cue-reward association events increase stimulus preference, indicating that such uncued rewards strengthen cue-reward associations.

Discussion

We monitored fMRI activity in visual cortex during uncued rewards that were separated in time from randomly interleaved cue-reward association trials. Surprisingly, fMRI activity monitored during these trials selectively decreased within the representation of the reward-predicting cue in visual cortex. Representation-specific decreases in fMRI activity were also found during the cue-reward association trials. These modulations were of smaller magnitude than uncued reward modulations supporting the hypothesis that the negative modulations we observed were dependent on PE. The similarity of the reward modulations during both cued and uncued trials, in conjunction with the dependence of uncued reward modulations on the presence of the cue-reward association, suggests that the online interaction of stimulus- and reward activity render reward modulations selective. The specificity of the uncued reward modulations was shown by the correlation of uncued reward- and cue-induced activity in experiments 1 and 3 and the ability to classify (via MVPA) the highly-rewarded cue by uncued reward activity in experiment 4. In addition, we found that the reduced fMRI activity observed during uncued rewards was dependent on several parameters: the size of uncued rewards, cue-reward probabilities, and cue-reward familiarity. Importantly, all of these effects could be explained by changes in the PE response during either the cue-reward association or the uncued reward. Furthermore, we found that selective reward modulations in visual cortex depended on dopamine signaling, as established through pharmacological intervention. Lastly, the uncued reward trials were found to strengthen behavioral cue-reward associations. These results are the first to show a cue-selective, negative, dopaminergic reward-feedback fMRI signal in visual cortex.

Selective Reduction in fMRI activity

We found that rewards *reduced* activity within the representation of reward-predicting cues during both cue-reward associations and uncued rewards. This is in contrast to previous studies, which have found either a lack of reward modulation (Weil et al., 2010) or increased activity for stimuli presented with rewards (Serences, 2008) within retinotopic visual cortex.

The stark differences found between studies likely results from critical differences in the experimental designs such as the inclusion of uncued reward trials in our study. Indeed, as shown in experiment 7, these uncued rewards clearly affect associations formed during cued-reward trials. In agreement with this, unpublished human experiments employing a similar design (i.e. with intermixed cue-reward and reward-only trials) have also revealed negative fMRI responses in visual cortex (T. Knapen, P. Roelfsema, J. Arsenault, W. Vanduffel, and T. Donner, personal communication, December 30, 2012).

Despite its robustness, negative reward activity is counterintuitive as one might expect a reward-predicting stimulus to be better-represented, and hence evoking increased activity. Yet the selective reduction in activity we observed may result in an enhanced representation of rewarded stimuli, a mechanism that may function more efficiently than increasing activity. For instance, the reduction in fMRI activity constitutes a dynamic (i.e. at the moment of reward delivery) and selective decrease in baseline activity within the cue-representation that subsequently boosts the signal-to-noise ratio during future cue presentations. Additionally, reward-induced deactivations may represent a decrease in overall activity with a simultaneous increase in stimulus information (Adab and Vogels, 2011; Kok et al., 2012). This is corroborated by Zalvidar et al. (2011), who found that visually-evoked fMRI activity was reduced by high doses of dopamine agonists. This decrease in fMRI activity was coupled with a concurrent increase in the signal-to-noise ratio for the stimulus. Thus, sparser coding of stimuli may be a highly efficient mechanism to enhance the representation of important stimuli, like those that predict reward.

Separating the Effects of Reward From Reward-related Cognitive Factors

One obstacle to interpreting the effects of reward associations on activity in sensory processing regions is the inherent difficulty of distinguishing reward from attentional effects, because attention is biased towards reward-predicting stimuli (Anderson et al., 2011; Peck et al., 2009). Therefore, while studies have found modulations within the visual representations of rewarded stimuli (Krawczyk et al., 2007; Serences, 2008) these effects were measured during stimulus presentation and discrimination, precisely when attentional bias is most likely to exist. Consequently, these studies cannot differentiate between the effects of attention and reward. In an effort to isolate such effects, other studies have temporally separated visual cue presentation from reward administration (Weil et al., 2010). Yet, in contrast to our work, these authors failed to find cue-specific reward modulation of fMRI activity in the retinotopic visual areas, although they did find an *interaction between attention and reward* within V3. The present report therefore demonstrates the first unambiguous evidence for a stimulus-selective reward signal in primate visual cortex. Furthermore, in contrast to the selective enhancements that have been observed within attended stimulus representations without visual stimulation (Kastner et al., 1999; Sylvester et al., 2007), we found a selective reduction of activity within the reward-paired cue representation. The opposite polarity of the reward modulations provides further evidence that the modulations we observed are unlikely to result from attention.

Hemodynamic activity in early visual cortex can display fluctuations that depend on trial structure and not reward (Sirotnin and Das, 2009), or upon the timing of the expected reward, rather than the reward itself (Shuler and Bear, 2006, their Figure 4). In experiment 1, uncued reward activity was defined by contrasting uncued reward trials with fixation trials. Crucially, the uncued rewards indicated the end of the current trial and the beginning of the next randomized wait period, while no information about trial structure was available during fixation trials. Trial-structure-dependent fluctuations in attention, hazard-rate or anticipation could therefore account for reward modulations observed in the first experiment. Alternatively, fixation trials in which no reward is administered could be viewed by the monkey as a reward-omission trial, leaving a reward-omission signal as a potential source of

the modulations recorded in experiment 1. To disambiguate this first set of results, we utilized a paradigm with two reward sizes, which conveyed the same trial structure information, in experiment 3. With trial-structure information held constant and reward omissions eliminated, we found significantly stronger deactivations within the cue-representation during larger uncued rewards. These results confirm that uncued reward activity was dependent on the attributes of the reward and not on other factors such as reward-omission or trial-structure.

A Dopaminergic Prediction Error Signal?

Manipulations of uncued reward size, cue-reward probabilities, and cue-reward familiarity have been shown to alter PE in monkeys and the subsequent responses of dopamine neurons (Schultz, 2006). For instance, large unpredicted rewards have been shown to elicit stronger PE and larger PE responses from dopamine neurons than small rewards (Tobler et al., 2005), exactly as we observed in the ventral midbrain (Experiment 3). Therefore, although we did not measure the monkey's subjective predictions directly through anticipatory licking (Fiorillo et al., 2003), the use of known properties of PE and the responses of dopamine neurons provided a consistent description of the data acquired in all 7 experiments. We would also like to note that while aspects of motivational functions controlled by dopamine can be accounted for by PE, PE obviously does not explain all dopaminergic functions in this complex domain (Salamone and Correa, 2012). Nonetheless, PE remains a useful construct when describing dopamine activity relative to transient changes in value.

Selective Reward Modulations Temporally-Surrounding Cue-Reward Associations

An important distinction between our experiments and prior studies that also separated stimulus presentation from reward (Pleger et al., 2008; Pleger et al., 2009; Weil et al., 2010) is that we measured modulations during rewards that were not part of discrete cue-reward association events. Hence, the reward modulations we observed in visual cortex demonstrate that events outside the actual cue-reward associations can selectively affect the representation of the reward-associated cue. This suggests, in conjunction with the reliance of uncued reward modulations on both the presence of cued trials (experiment 2) and properties of the cue-reward association (experiment 4 and 5), that the degree and location of uncued reward modulations is controlled by a two-stage process during cue-reward and uncued reward trials, respectively. We hypothesize that the interaction of cue-specific sensory activity and a more diffuse reward-driven feedback signal “tag” the stimulus representation. Thereafter, a diffuse reward signal is generated by the uncued reward that preferentially interacts with the previously “tagged” stimulus representation, creating a selective reward modulation at the cue-representation.

The increase in the monkey's cue preference monitored when cue-reward association trials were surrounded by uncued rewards (experiment 7) provides further evidence for a two-stage process in which uncued rewards affect the associations formed during cue-reward trials. Furthermore, this effect strongly refutes the hypothesis that uncued rewards and the modulations we observed represent a weakening of the cue-reward relationship. Additional studies must be conducted to determine whether factors like uncued reward probability and the timing of rewards strengthen or weaken cue-reward relationships. More generally, the strengthening of the reward-association that we monitored is in agreement with a body of work showing that dopamine-releasing events, temporally separated from learning events, facilitate learning (White and Milner, 1992; Wise, 2004). The specificity of these behavioral enhancements to the learned event suggests that the widespread dopamine signal is somehow rendered selective to the representation of the learned event. It is therefore tempting to speculate that the cue-selective dopamine-dependent signal we have shown may represent a general mechanism through which dopamine signals become selective.

Potential Sources of Reward Modulation in Visual Cortex

Manipulations of both the cue-reward association (experiment 2, 4 & 5) and the uncued reward (experiment 3) indicate that PE during these events determines the strength and location of uncued reward modulations. The influence of perturbations in PE on uncued reward activity, in conjunction with its susceptibility to dopamine antagonist application (experiment 6), indicates that uncued reward activity may be regulated by a dopaminergic PE signal, potentially originating in the ventral midbrain. This is further supported by the run-by-run correlation in experiment 3 between activity within the cue-representation and the ventral midbrain during uncued rewards. This evidence suggests that the observed activity modulations in visual cortex are indeed caused by a dopaminergic PE signal. An important question remaining is whether the spatially selective effects are induced by the specificity of top-down or bottom-up projections to visual cortex that can be functionally modulated by dopamine (Noudoost and Moore, 2011; Zhao et al., 2002), or, alternatively, result from sparser dopaminergic connections between ventral midbrain and visual cortex.

EXPERIMENTAL PROCEDURES

Subjects

All procedures were approved by the KUL's Committee on Animal Care, and are in accordance with NIH and European guidelines for the care and use of laboratory animals. Eight rhesus monkeys (*Macaca mulatta*; M13, M18, M19, M20, M22, M23, M26, M9 4.5-7 kg, 6-9 years old, 7 males) were trained for a passive fixation task and prepared for awake fMRI as previously describe (Vanduffel et al., 2001). For the two monkeys (M19, M20) that participated in the pharmacological challenge experiment, a catheter (silicone; 0.7mm inner diameter; Access Technologies) was chronically inserted into the internal jugular vein see (Nelissen et al., 2012), see Supplemental Experimental Procedures)

Functional MRI Acquisition

Contrast-agent-enhanced functional images (Leite et al., 2002; Vanduffel et al., 2001) were acquired in a 3.0 T horizontal bore full-body scanner (TIM Trio, Siemens Healthcare; Erlangen, Germany), using a gradient-echo T2* weighted echo-planar sequence (50 horizontal slices, in-plane 84 × 84 matrix, TR = 2s, TE = 19 ms, 1 × 1 × 1 mm³ isotropic voxels). An eight-channel phased array coil system (individual coils 3.5 cm diameter), with offline SENSE reconstruction, an image acceleration factor of 3, and a saddle-shaped, radial transmit-only surface coil were employed (Kolster et al., 2009).

Cue Localizer Experiment (N = 4; M13, M18, M19, M20)

fMRI responses to the abstract visual stimuli (red and green cues, see Figure S1A and S1B) presented for 500 ms with a 3500 – 6000 ms inter-stimulus interval were measured during independent localizer scans (see Supplemental Experimental Procedures). The form of the visual stimuli was similar to stimuli used in a previous experiment (Pessiglione et al., 2006). Note that within this localizer experiment, the visual stimuli did not predict upcoming reward. This goal was achieved by presenting the reward and the stimulus events on asynchronous time schedules. Three equiprobable events (green cue, red cue and fixation) occurred every 3500 - 6000 ms (actual inter-stimulus intervals were generated randomly on each run) and lasted for 500 ms while juice rewards were administered every ~1000 ms. Statistical thresholds, the number of runs used to define the cue-representation ROIs, and the figures for which data from a given ROI was used are displayed in Table S1.

Experiment 1: 2-by-2 Factorial Design (N = 3; M13, M18, M19)

This design consisted of four equiprobable trial types (fixation, uncued reward, cue, and cue-reward). The monkeys had to maintain fixation within a $2 \times 3^\circ$ window during a randomly jittered 3.5 – 6s waiting period. During cue-reward trials, a ~6-deg abstract green line drawing (see Figure S1A) appeared for 500 ms, and 400 ms after cue onset a 0.2 ml juice reward was administered (cue-reward). The timing of the visual cue and the reward was held constant in the cue and uncued reward trials, respectively. During a fixation trial, no visual stimulus was presented but a 500 ms window was added to keep the trial duration the same.

Experiment 2: Uncued Rewards Without Temporally-Surrounding Cued Trials Types (N=2; M22, M23)

This design was identical to experiment 1 although all cued trial types were omitted (cue and cue-reward). Therefore experiment 2 consisted solely of fixation and uncued reward trials. The animals that performed this experiment were never exposed to the direct pairing of the juice reward and the visual cues.

Experiment 3: Small and Large Uncued Reward Level Experiment (N = 2; M19, M20)

The reward-level experiment was identical to experiment 1 except that it consisted of both a small (0.1 ml) and a large (0.3ml) uncued reward condition rather than the single uncued reward condition (0.2 ml).

Experiment 4: Cue-reward Probability Experiment (N = 2; M19, M20)

In this experiment there were 3 condition groups [green cue (Figure S1A), red cue (Figure S1B), and uncued], all of which were equiprobable. There were two variants of this experiment (green and red high reward-probability). During the green high reward-probability experiments, the green cue was followed by rewards in 66% of the trials while the red cue was followed by reward in 33% of the trials. For red high reward-probability experiments, the cue-reward probabilities were reversed. During both green and red high-reward experiments, uncued trials were rewarded 50% of the time. In addition, the order of the green and red high-reward experiments was counterbalanced between subjects (Figure 6C - 6D).

Experiment 5: Cue-reward Familiarity Experiment (N = 2; M19, M20)

This paradigm was identical to experiment 4, with the exception that one of the cues was invariably followed by a reward (100% of trials rewarded, M19 - green cue, M20 - red cue) while the other cue was never rewarded. Significantly, before this experiment began, monkeys were trained in a paradigm where both the green and red cues were rewarded 50% of the time (number of training runs, M19 – 50 runs, M20 – 41 runs).

Experiment 6: Dopamine Antagonist (SCH-23390) Experiment (N = 2; M19, M20)

The experimental paradigm was identical to experiment 1 (runs consisted of equiprobable fixation, uncued reward, cue and cue-reward trial types) with the exception that during one of the runs, two boluses of a D1-selective dopamine antagonist were injected. Experimental sessions were separated into baseline (immediately preceding the injection run), post-injection (immediately following the injection run) and recovery (directly following the post-injection runs) phases. The 3 phases were equalized for scan time (3 runs/phase, 305 volumes/run, 2 s/volume) and number of events per condition (baseline, post-injection, recovery; M19, 84.9 events/condition/phase; M20, 82.8 events/condition/phase). The injection run was excluded from fMRI analysis but consisted of 2 small bolus injections (duration 30 s) via a jugular catheter, of (0.0025 - 0.005 mg/kg) selective D1 antagonist

R(+)-SCH-23390 hydrochloride (Sigma-Aldrich; St Louis, MO) five minutes apart. In any given session, both injections were of the same concentration. Two injections were administered rather than a single dose to limit potential extrapyramidal effects associated with peak concentrations of dopamine antagonist (Fischer et al., 2010). Each animal participated in 5 sessions, resulting in 15 runs/phase/animal and 30 runs/phase in total. Injection of SCH-23390 into rats has been shown to have a 30 min half-life in plasma while displaying a slightly longer half-life of 40–60 min in the striatum and cortex (Hietala et al., 1992). Therefore, the runs following the post-injection phase were deemed the recovery phase with the caveat that physiological relevant levels of SCH-23390 may still be present in the brain, albeit at a lower concentration than in the post-injection phase.

Experiment 7: Behavioral Effects of Uncued Rewards (N = 2; M9, M26)

Each session contained a cue-reward association block and two free-choice stimulus preference tests (400 trials). The first preference test preceded the association block while the second test immediately followed it. Preference tests were used to assess potential changes in stimulus preference. Stimulus preference trials began when the animal fixated on a central fixation point. After 1000–1500 ms, two stimuli (~7-deg in size) were simultaneously presented peripheral (9.5-deg eccentricity) to the fixation point for up to 2000 ms, one to the left and the other to the right of the fixation point. For each session, 2 novel stimuli were chosen from a randomized set of basic geometric shapes that differed in both shape and color. A trial was completed and the stimuli were removed after a saccade to one of the two stimuli. The position of the stimuli was randomly alternated and both stimuli were rewarded with a 50% reward probability. After testing stimulus preference, the less-selected stimulus (i.e. non-preferred stimulus) was associated with a juice reward during 25 cue-reward trials within a cue-reward association block. There were two variants of the cue-reward association blocks, those that contained uncued reward trials and those that did not. Association blocks with uncued rewards were identical to experiment 1 and therefore contained 4 equiprobable trial types (fixation, reward, cue, cue-reward). Association blocks without uncued rewards contained 2 equiprobable trial types (cue and cue-reward). After the cue-reward association block another stimulus preference test was performed. Analysis was performed in 20 trial bins comparing non-preferred stimulus selection before and after the two different types of cue-reward association blocks.

General Linear Model (GLM) Analysis

Images were first reconstructed then realigned using a non-rigid slice-by-slice registration algorithm (Kolster et al., 2009). The resultant images were next 3D motion-corrected within session, smoothed (FWHM 1.5 mm), and non-rigidly co-registered to each subject's own anatomical template using Match Software (Chef d'Hotel et al., 2002). We then performed a voxel-based analysis of with SPM5, following previously-described procedures to fit a general linear model (Friston et al., 1995; Leite et al., 2002; Vanduffel et al., 2001; Vanduffel et al., 2002). High- and low-pass filtering were employed prior to fitting the GLM. To account for head- and eye-movement related artifacts, 6 motion-realignment parameters and 2 eye parameters were used as covariates of no interest. Eye traces were thresholded within the $2^\circ \times 3^\circ$ window, convolved with the MION response function and subsampled to the TR (2s).

Region of Interest (ROI) Definition

The borders of 6 visual areas (V1, V2, V3, V4, TEO, and TE) were identified on a flattened cortical representation (Van Essen et al., 2001) using retinotopic mapping data previously collected in three animals (Fize et al., 2003) and an atlas (Ungerleider and Desimone, 1986) co-registered to the flattened cortical representation. To define the cue-representations, we determined the subset of voxels, within each visual area, that were activated during the

localizer experiment (see Table S1). Midbrain functional ROIs were defined as midbrain voxels maximally driven by uncued rewards [5 mm³ each hemisphere; (small uncued reward + large uncued reward) – fixation; M19, $T > 5.2$; M20, $T > 10.6$]. In addition, we non-linearly transformed our midbrain ROIs into an atlas space (Saleem and Logothetis, 2006) and confirmed their colocalization with the ventral tegmental area.

Eye Position Analysis (Experiments 1-6)

Eye position was continuously monitored with an infrared pupil/corneal reflection tracking system (120 Hz) over a 10-second window surrounding cue presentation (4 seconds before cue onset to 6 seconds after). Percent fixation within the 2-by-3 degree window of eye position was compared between conditions for this time window. Either a Wilcoxon rank sum test or a Kruska-Wallis non-parametric ANOVA was used to calculate significances of differences between conditions (see Tables S2-7).

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

We thank C. Fransen, C. Van Eupen and A. Coeman for animal training and care; D. Mantini, O. Joly, H. Kolster, W. Depuydt, G. Meulemans, P. Kayenbergh, M. De Paep, M. Docx, and I. Puttemans for technical assistance; and P. Roelfsema, T. Knapen, T. Donner and S. Raiguel for their comments on the manuscript. This work received support from Inter-University Attraction Pole 7/21, Programme Financing PFV/10/008, Geconcerteerde Onderzoeks Actie 10/19, Impulsfinanciering Zware Apparatuur and Hercules funding of the Katholieke Universiteit Leuven, Fonds Wetenschappelijk Onderzoek–Vlaanderen G062208.10, G083111.10 and G.0719.12, and G0888.13. K.N. is postdoctoral fellow of the Fonds Wetenschappelijk Onderzoek–Vlaanderen. The Martinos Center for Biomedical Imaging is supported by National Center for Research Resources grant P41RR14075.

REFERENCES

- Adab HZ, Vogels R. Practicing coarse orientation discrimination improves orientation signals in macaque cortical area v4. *Curr Biol*. 2011; 21:1661–1666. [PubMed: 21962714]
- Anderson BA, Laurent PA, Yantis S. Value-driven attentional capture. *Proc Natl Acad Sci U S A*. 2011; 108:10367–10371. [PubMed: 21646524]
- Bao S, Chan VT, Merzenich MM. Cortical remodelling induced by activity of ventral tegmental dopamine neurons. *Nature*. 2001; 412:79–83. [PubMed: 11452310]
- Blaiss CA, Janak PH. Post-training, but not post-reactivation, administration of amphetamine and anisomycin modulates Pavlovian conditioned approach. *Neurobiol Learn Mem*. 2007; 87:644–658. [PubMed: 17258476]
- Bromberg-Martin ES, Matsumoto M, Hikosaka O. Dopamine in motivational control: rewarding, aversive, and alerting. *Neuron*. 2010; 68:815–834. [PubMed: 21144997]
- Chef d'Hotel C, Hermsillo G, Faugeras O. Flows of diffeomorphisms for multimodal image registration. *Proc IEEE Int S Bio Im*. 2002; 7-8:753–756.
- Engelmann JB, Damaraju E, Padmala S, Pessoa L. Combined effects of attention and motivation on visual task performance: transient and sustained motivational effects. *Front Hum Neurosci*. 2009; 3:4. [PubMed: 19434242]
- Engelmann JB, Pessoa L. Motivation sharpens exogenous spatial attention. *Emotion*. 2007; 7:668–674. [PubMed: 17683222]
- Fiorillo CD, Tobler PN, Schultz W. Discrete coding of reward probability and uncertainty by dopamine neurons. *Science*. 2003; 299:1898–1902. [PubMed: 12649484]
- Fischer H, Nyberg L, Karlsson S, Karlsson P, Brehmer Y, Rieckmann A, MacDonald SW, Farde L, Backman L. Simulating neurocognitive aging: effects of a dopaminergic antagonist on brain activity during working memory. *Biol Psychiatry*. 2010; 67:575–580. [PubMed: 20138255]

- Fize D, Vanduffel W, Nelissen K, Denys K, Chef d'Hotel C, Faugeras O, Orban GA. The retinotopic organization of primate dorsal V4 and surrounding areas: A functional magnetic resonance imaging study in awake monkeys. *J Neurosci*. 2003; 23:7395–7406. [PubMed: 12917375]
- Friston KJ, Holmes AP, Poline JB, Grasby PJ, Williams SC, Frackowiak RS, Turner R. Analysis of fMRI time-series revisited. *Neuroimage*. 1995; 2:45–53. [PubMed: 9343589]
- Gonon FG. Nonlinear relationship between impulse flow and dopamine released by rat midbrain dopaminergic neurons as studied by in vivo electrochemistry. *Neuroscience*. 1988; 24:19–28. [PubMed: 3368048]
- Hietala J, Seppala T, Lappalainen J, Syvalahti E. Quantification of SCH 39166, a novel selective D1 dopamine receptor antagonist, in rat brain and blood. *Psychopharmacology (Berl)*. 1992; 106:455–458. [PubMed: 1349751]
- Hitchcott PK, Phillips GD. Double dissociation of the behavioural effects of R(+) 7-OH-DPAT infusions in the central and basolateral amygdala nuclei upon Pavlovian and instrumental conditioned appetitive behaviours. *Psychopharmacology (Berl)*. 1998; 140:458–469. [PubMed: 9888622]
- Hollerman JR, Schultz W. Dopamine neurons report an error in the temporal prediction of reward during learning. *Nat Neurosci*. 1998; 1:304–309. [PubMed: 10195164]
- Kastner S, Pinsk MA, De Weerd P, Desimone R, Ungerleider LG. Increased activity in human visual cortex during directed attention in the absence of visual stimulation. *Neuron*. 1999; 22:751–761. [PubMed: 10230795]
- Kok P, Jehee JF, de Lange FP. Less is more: expectation sharpens representations in the primary visual cortex. *Neuron*. 2012; 75:265–270. [PubMed: 22841311]
- Kolster H, Mandeville JB, Arsenault JT, Ekstrom LB, Wald LL, Vanduffel W. Visual field map clusters in macaque extrastriate visual cortex. *J Neurosci*. 2009; 29:7031–7039. [PubMed: 19474330]
- Krawczyk DC, Gazzaley A, D'Esposito M. Reward modulation of prefrontal and visual association cortex during an incentive working memory task. *Brain Res*. 2007; 1141:168–177. [PubMed: 17320835]
- Leite FP, Tsao D, Vanduffel W, Fize D, Sasaki Y, Wald LL, Dale AM, Kwong KK, Orban GA, Rosen BR, et al. Repeated fMRI using iron oxide contrast agent in awake, behaving macaques at 3 Tesla. *Neuroimage*. 2002; 16:283–294. [PubMed: 12030817]
- Lidow MS, Goldman-Rakic PS, Gallager DW, Rakic P. Distribution of dopaminergic receptors in the primate cerebral cortex: quantitative autoradiographic analysis using [3H]raclopride, [3H]spiperone and [3H]SCH23390. *Neuroscience*. 1991; 40:657–671. [PubMed: 2062437]
- Lindvall O, Bjorklund A, Moore RY, Stenevi U. Mesencephalic dopamine neurons projecting to neocortex. *Brain Res*. 1974; 81:325–331. [PubMed: 4373129]
- Ljungberg T, Apicella P, Schultz W. Responses of monkey dopamine neurons during learning of behavioral reactions. *J Neurophysiol*. 1992; 67:145–163. [PubMed: 1552316]
- Matsumoto M, Hikosaka O. Two types of dopamine neuron distinctly convey positive and negative motivational signals. *Nature*. 2009; 459:837–841. [PubMed: 19448610]
- Maunsell JH. Neuronal representations of cognitive state: reward or attention? *Trends Cogn Sci*. 2004; 8:261–265. [PubMed: 15165551]
- Messier C, White NM. Contingent and non-contingent actions of sucrose and saccharin reinforcers: effects on taste preference and memory. *Physiol Behav*. 1984; 32:195–203. [PubMed: 6718546]
- Nelissen K, Jarraya B, Arsenault JT, Rosen BR, Wald LL, Mandeville JB, Marota JJ, Vanduffel W. Neural Correlates of the Formation and Retention of Cocaine-Induced Stimulus-Reward Associations. *Biol Psychiatry*. 2012
- Nomoto K, Schultz W, Watanabe T, Sakagami M. Temporally extended dopamine responses to perceptually demanding reward-predictive stimuli. *J Neurosci*. 2010; 30:10692–10702. [PubMed: 20702700]
- Noudoost B, Moore T. Control of visual cortical signals by prefrontal dopamine. *Nature*. 2011; 474:372–375. [PubMed: 21572439]
- O'Doherty J, Dayan P, Schultz J, Deichmann R, Friston K, Dolan RJ. Dissociable roles of ventral and dorsal striatum in instrumental conditioning. *Science*. 2004; 304:452–454. [PubMed: 15087550]

- Oscos A, Martinez JL Jr, McGaugh JL. Effects of post-training d-amphetamine on acquisition of an appetitive autoshaped lever press response in rats. *Psychopharmacology (Berl)*. 1988; 95:132–134. [PubMed: 3133692]
- Peck CJ, Jangraw DC, Suzuki M, Efem R, Gottlieb J. Reward modulates attention independently of action value in posterior parietal cortex. *J Neurosci*. 2009; 29:11182–11191. [PubMed: 19741125]
- Pessiglione M, Petrovic P, Daunizeau J, Palminteri S, Dolan RJ, Frith CD. Subliminal instrumental conditioning demonstrated in the human brain. *Neuron*. 2008; 59:561–567. [PubMed: 18760693]
- Pessiglione M, Seymour B, Flandin G, Dolan RJ, Frith CD. Dopamine-dependent prediction errors underpin reward-seeking behaviour in humans. *Nature*. 2006; 442:1042–1045. [PubMed: 16929307]
- Pleger B, Blankenburg F, Ruff CC, Driver J, Dolan RJ. Reward facilitates tactile judgments and modulates hemodynamic responses in human primary somatosensory cortex. *J Neurosci*. 2008; 28:8161–8168. [PubMed: 18701678]
- Pleger B, Ruff CC, Blankenburg F, Kloppel S, Driver J, Dolan RJ. Influence of dopaminergically mediated reward on somatosensory decision-making. *PLoS Biol*. 2009; 7:e1000164. [PubMed: 19636360]
- Roelfsema PR, van Ooyen A, Watanabe T. Perceptual learning rules based on reinforcers and attention. *Trends Cogn Sci*. 2010; 14:64–71. [PubMed: 20060771]
- Roesch MR, Olson CR. Neuronal activity related to reward value and motivation in primate frontal cortex. *Science*. 2004; 304:307–310. [PubMed: 15073380]
- Salamone JD, Correa M. The mysterious motivational functions of mesolimbic dopamine. *Neuron*. 2012; 76:470–485. [PubMed: 23141060]
- Saleem, KS.; Logothetis, NK. *Combined MRI and Histology Atlas of the Rhesus Monkey Brain*. Academic Press; Amsterdam: 2006.
- Schultz W. Behavioral theories and the neurophysiology of reward. *Annu Rev Psychol*. 2006; 57:87–115. [PubMed: 16318590]
- Schultz W, Dayan P, Montague PR. A neural substrate of prediction and reward. *Science*. 1997; 275:1593–1599. [PubMed: 9054347]
- Seitz A, Watanabe T. A unified model for perceptual learning. *Trends Cogn Sci*. 2005; 9:329–334. [PubMed: 15955722]
- Seitz AR, Kim D, Watanabe T. Rewards evoke learning of unconsciously processed visual stimuli in adult humans. *Neuron*. 2009; 61:700–707. [PubMed: 19285467]
- Serences JT. Value-based modulations in human visual cortex. *Neuron*. 2008; 60:1169–1181. [PubMed: 19109919]
- Shuler MG, Bear MF. Reward timing in the primary visual cortex. *Science*. 2006; 311:1606–1609. [PubMed: 16543459]
- Sirotin YB, Das A. Anticipatory haemodynamic signals in sensory cortex not predicted by local neuronal activity. *Nature*. 2009; 457:475–479. [PubMed: 19158795]
- Surmeier DJ, Shen W, Day M, Gertler T, Chan S, Tian X, Plotkin JL. The role of dopamine in modulating the structure and function of striatal circuits. *Prog Brain Res*. 2010; 183:149–167. [PubMed: 20696319]
- Sylvester CM, Shulman GL, Jack AI, Corbetta M. Asymmetry of anticipatory activity in visual cortex predicts the locus of attention and perception. *J Neurosci*. 2007; 27:14424–14433. [PubMed: 18160650]
- Tan CO. Anticipatory changes in regional cerebral hemodynamics: a new role for dopamine? *J Neurophysiol*. 2009; 101:2738–2740. [PubMed: 19321643]
- Tobler PN, Fiorillo CD, Schultz W. Adaptive coding of reward value by dopamine neurons. *Science*. 2005; 307:1642–1645. [PubMed: 15761155]
- Ungerleider LG, Desimone R. Projections to the superior temporal sulcus from the central and peripheral field representations of V1 and V2. *J Comp Neurol*. 1986; 248:147–163. [PubMed: 3722456]

- Van Essen DC, Lewis JW, Drury HA, Hadjikhani N, Tootell RB, Bakircioglu M, Miller MI. Mapping visual cortex in monkeys and humans using surface-based atlases. *Vision Res.* 2001; 41:1359–1378. [PubMed: 11322980]
- Vanduffel W, Fize D, Mandeville JB, Nelissen K, Van Hecke P, Rosen BR, Tootell RB, Orban GA. Visual motion processing investigated using contrast agent-enhanced fMRI in awake behaving monkeys. *Neuron.* 2001; 32:565–577. [PubMed: 11719199]
- Vanduffel W, Fize D, Peuskens H, Denys K, Sunaert S, Todd JT, Orban GA. Extracting 3D from motion: differences in human and monkey intraparietal cortex. *Science.* 2002; 298:413–415. [PubMed: 12376701]
- Vickery TJ, Chun MM, Lee D. Ubiquity and specificity of reinforcement signals throughout the human brain. *Neuron.* 2011; 72:166–177. [PubMed: 21982377]
- Wang SH, Redondo RL, Morris RG. Relevance of synaptic tagging and capture to the persistence of long-term potentiation and everyday spatial memory. *Proc Natl Acad Sci U S A.* 2010; 107:19537–19542. [PubMed: 20962282]
- Weil RS, Furl N, Ruff CC, Symmonds M, Flandin G, Dolan RJ, Driver J, Rees G. Rewarding feedback after correct visual discriminations has both general and specific influences on visual cortex. *J Neurophysiol.* 2010; 104:1746–1757. [PubMed: 20660419]
- White N, Major R. Effect of pimozide on the improvement in learning produced by self-stimulation and by water reinforcement. *Pharmacol Biochem Behav.* 1978; 8:565–571. [PubMed: 674262]
- White NM, Milner PM. The psychobiology of reinforcers. *Annu Rev Psychol.* 1992; 43:443–471. [PubMed: 1539948]
- Wise RA. Dopamine, learning and motivation. *Nat Rev Neurosci.* 2004; 5:483–494. [PubMed: 15152198]
- Zalvidar, D.; Von Pfoestl, JLV.; Zhang, X.; Logothetis, N.; Rauch, A. Soc Neurosci Abstr. Washington D.C.: 2011. Dopaminergic modulation of the early visual system of non-human primates and its underlying neuronal and hemodynamic changes.
- Zhang L, Doyon WM, Clark JJ, Phillips PE, Dani JA. Controls of tonic and phasic dopamine transmission in the dorsal and ventral striatum. *Mol Pharmacol.* 2009; 76:396–404. [PubMed: 19460877]
- Zhao Y, Kerscher N, Eysel U, Funke K. D1 and D2 receptor-mediated dopaminergic modulation of visual responses in cat dorsal lateral geniculate nucleus. *J Physiol.* 2002; 539:223–238. [PubMed: 11850515]

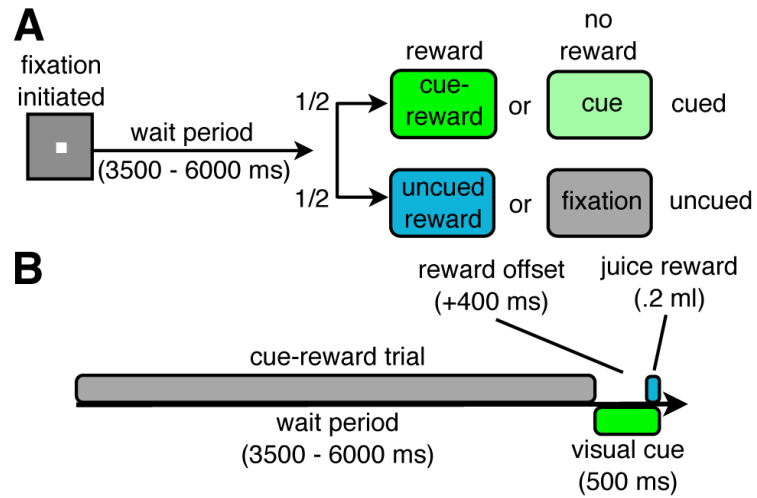


Figure 1. Design of 2-by-2 experiment (experiment 1)

(A) 2-by-2 factorial design (visual cue and juice reward). All event types (fixation, uncued reward, cue, cue-reward) were equiprobable. (B) Timing of an individual cue-reward trial. The timing of the juice reward and the visual cue were the same in the uncued reward and cue trials, respectively. See also Table S1 and Figure S1.

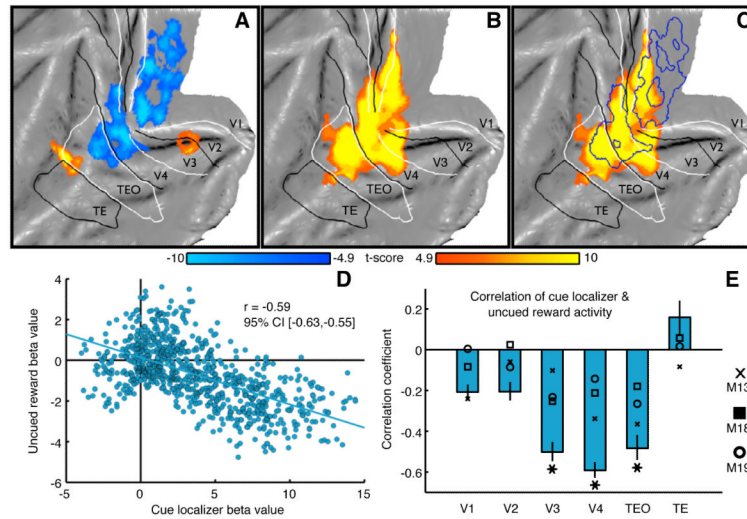


Figure 2. Spatially-selective reward-induced deactivations in visual cortex without visual stimulation (experiment 1)

(A) Uncued reward fMRI activity measured in 2-by-2 experiment (uncued reward – fixation, $P < 0.05$ FWE corrected, group-level analysis, 181 runs, M13 - 61 runs, M18 - 62 runs, M19 - 58 runs) and (B) cue localizer fMRI activity (see Table S1) projected onto a flattened representation of left occipital cortex. (C) Uncued reward activity (blue outline) overlaid onto cue localizer activity (orange-yellow). (D) Beta values from group-level analysis of all voxels in area V4 (871 voxels) for visual cue (x-axis, green cue - fixation, localizer experiment) and uncued rewards (y-axis, uncued reward - fixation, experiment 1). (E) Correlation of cue localizer and uncued reward activity across all voxels in a given visual region (group-level analysis). Blue bars represent the mean correlation coefficients. Error bars represent the 95% confidence interval (CI) of the correlation coefficients computed using a bootstrap algorithm (5000 samples). Symbols denote the mean correlation coefficients of separate single-subject analyses [M13 (cross); M18 (square); M19 (circle)]. Asterisks denote significant correlations in the group and in all individuals ($P < 0.05$, Bonferroni–Holm corrected for multiple comparisons across 6 ROIs). See also Table S2 and Figure S2.

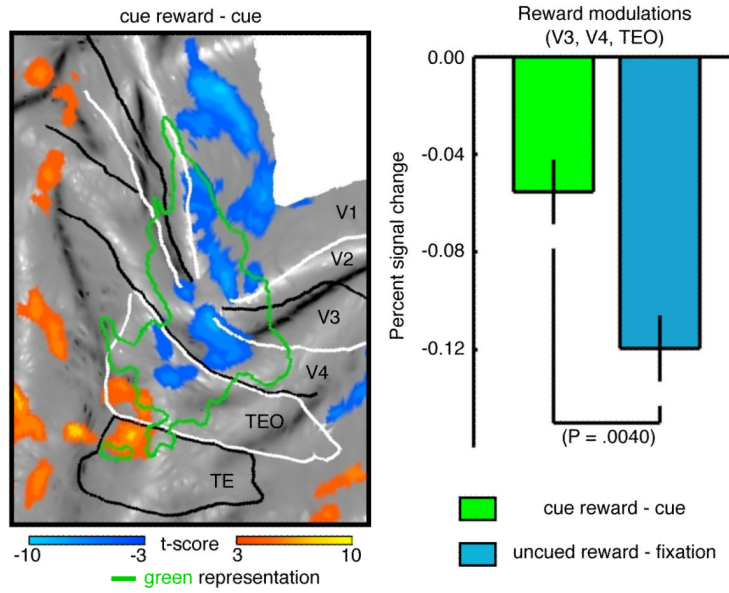


Figure 3. Comparison of reward modulations during cued & uncued trials

(A) Reward modulation during cued trials projected onto a flattened representation of left occipital cortex (cue reward – cue, $P < 0.001$, group-level analysis, 181 runs, M13 - 61 runs, M18 - 62 runs, M19 - 58 runs) and the cue-representation (see Table S1). (B) Group mean reward modulation PSC within the cue-representation (V3, V4 and TEO, see Table S1) during cued (cue reward – cue) and uncued trials (uncued reward – fixation). Error bars denote the SEM across runs. Significance determined using Wilcoxon rank sum test.

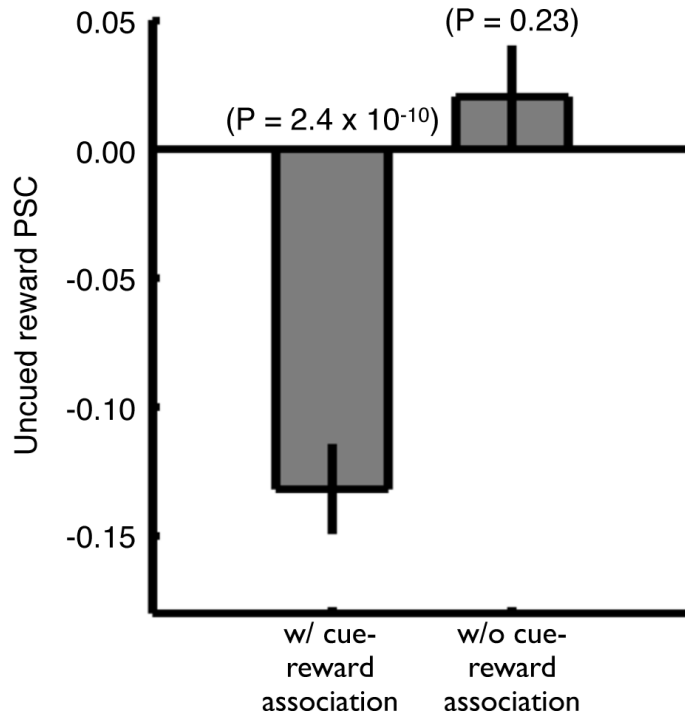


Figure 4. Uncued reward activity with and without surrounding cue-reward trials (Experiments 1 & 2)

Group mean uncued reward PSC within the cue-representation (V3, V4 and TEO, see Table S1) during experiment 1 (with cue-reward association, 80 runs, M18 - 40 runs, M19 - 40 runs) and experiment 2 (without cue-reward association, 80 runs, M22 - 40 runs, M23 - 40 runs). Error bars denote the SEM across runs. Significance determined using a Wilcoxon signed rank test. See also Table S3 and Figure S3.

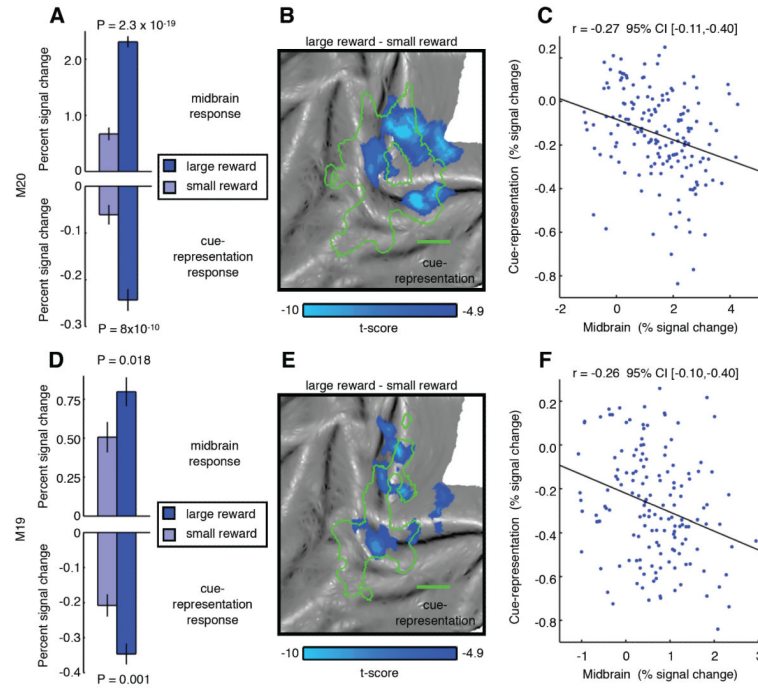


Figure 5. Size of uncued reward modulates fMRI activity within ventral midbrain and visual cue-representation (experiment 3)

Mean PSC of large reward and small uncued rewards relative to fixation trials for (A) M20 (82 runs) and (D) M19 (76 runs) measured within a ventral midbrain ROI (see Experimental Procedures) and the cue-representation (see Table S1). Error bars denote SEM across runs. Significance between large and small uncued rewards PSC determined using Wilcoxon rank sum test. (M20 - B, M19 - E). Difference in fMRI activity between large and small uncued rewards (large uncued reward - small uncued reward, $P < 0.05$ FWE corrected) projected onto a flattened cortical representation of left occipital cortex with cue localizer activity boundaries overlaid (green outline). (M20 - C, M19 - F) Run-by-run PSC for reward conditions (small uncued reward and large uncued rewards) relative to fixation trials monitored in the ventral midbrain and the cue-representation. Black line denotes the least squares line of best fit. Significance of correlation determined using bootstrap algorithm (5000 samples) to estimate the 95% CI. See also Table S4 and Figure S4.

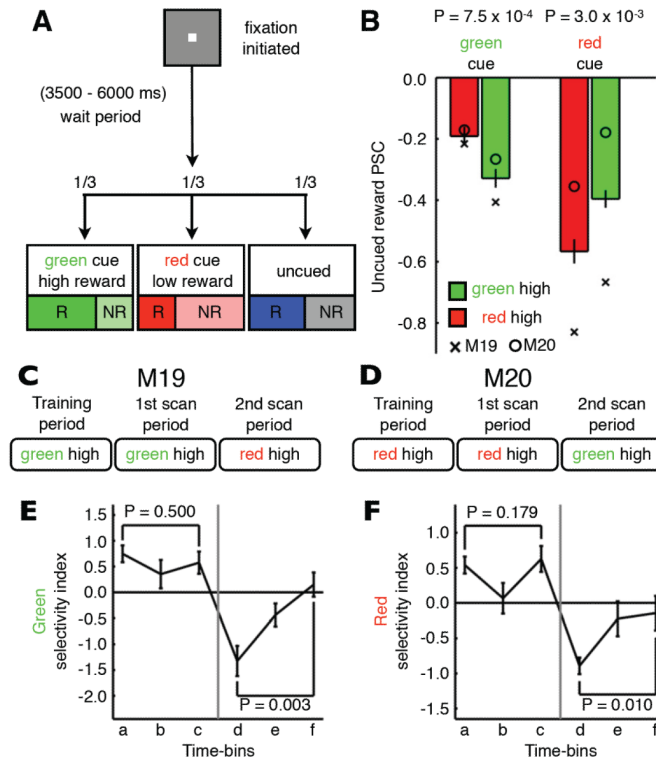


Figure 6. Cue-reward relationships modulate uncued reward activity in visual cortex
 (A) Design of cue-value experiment (experiment 4) (green, high reward-probability example, 66% green cue rewarded (R), 66% red cue not rewarded (NR), 50% of uncued trials rewarded). (B) Mean PSC during uncued rewards (uncued reward - fixation, 117 runs, M19 - 66 runs, M20 - 51 runs) measured within the green (left) and red (right) cue-representations (see Table S1). Bar color indicates the high reward-probability cue. Error bars denote SEM across runs. Symbols denote the mean PSC of individuals [M19 (cross); M20 (circle)]. Wilcoxon rank sum test compared PSC during green and red high reward-probability scan periods. The timeline of the cue-reward relationships that (C) M19 and (D) M20 were exposed to during the training phase and first and second scan periods. (E) Mean green-cue selectivity index over the timecourse of the experiment (M19, 2 scan periods, 3 time-bins/scan period, 22 runs/time-bin). (F) Mean red cue selectivity index over the timecourse of the experiment (M20, 2 scan periods, 3 time-bins/scan period, 17 runs/time-bin). In both (E) and (F), error bars denote SEM across runs and significance was determined using Kruskal-Wallis non-parametric ANOVA across the three time-bins in a given scan period. See also Table S5 and Figure S5.

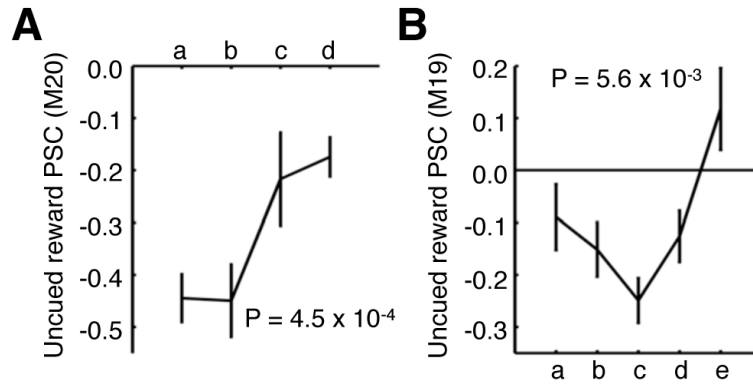


Figure 7. Uncued reward activity decreases after prolonged exposure to an absolute cue-reward relationship (experiment 5)

Mean PSC of uncued rewards (uncued reward – fixation) separated into equal-length time-bins within the 100% reward predicting cue-representation for (A) M20 (green representation, 15 runs/time-bin) and (B) M19 (red representation, 14 runs/time-bin). Time-bins comprise runs acquired at progressively later time points during the experiment. See Table S1 for cue-representation definitions. Error bars denote SEM. Significance determined using Kruskal-Wallis non-parametric ANOVA comparing the PSC across time-bins. See also Table S6 and Figure S6.

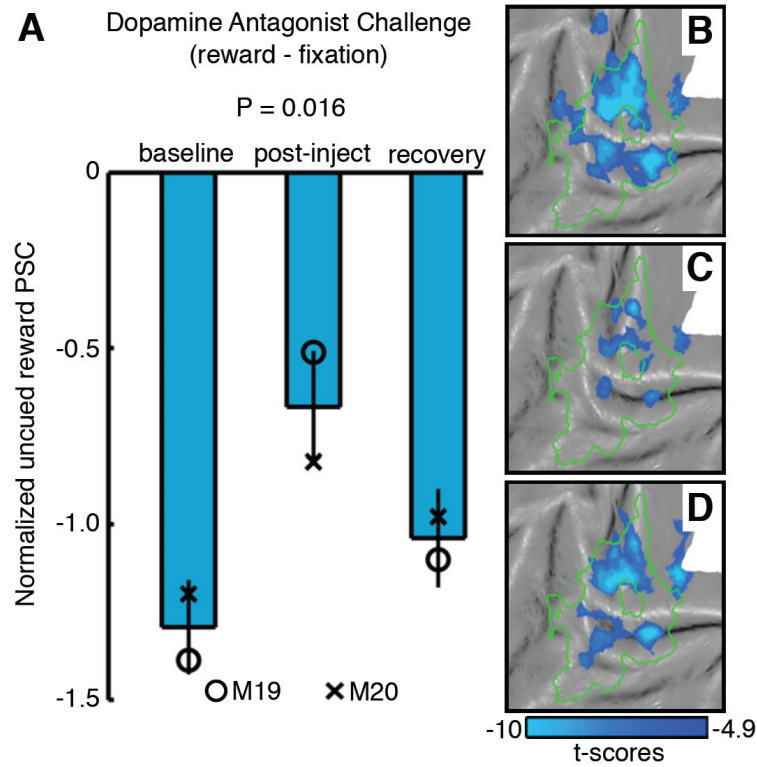


Figure 8. Uncued reward activity in visual cortex is susceptible to dopamine D1-receptor antagonist (SCH-23390) challenge (experiment 6)

(A) Mean normalized PSC (see Supplemental Experimental Procedures) during uncued rewards (uncued reward – fixation, group-level analysis, 30 runs/phase, M19 & M20 - 15 runs/phase) within the cue-representation (see Table S1) measured during baseline, post-injection and recovery phases. Error bars denote SEM across runs. Symbols denote the mean normalized PSC of single-subject analyses [M19 (circle); M20 (cross)]. Kruskal-Wallis non-parametric ANOVA was performed comparing PSC across phases. Uncued reward fMRI activity (uncued reward – fixation, $P < 0.05$ FWE corrected group-level analysis, 30 runs/phase, M19 & M20 - 15 runs/phase) projected onto a flattened cortical representation of left occipital cortex during the (B) baseline, (C) post-injection, and (D) recovery phases. Green outline represents the cue-representation. See also Table S7 and Figure S8.

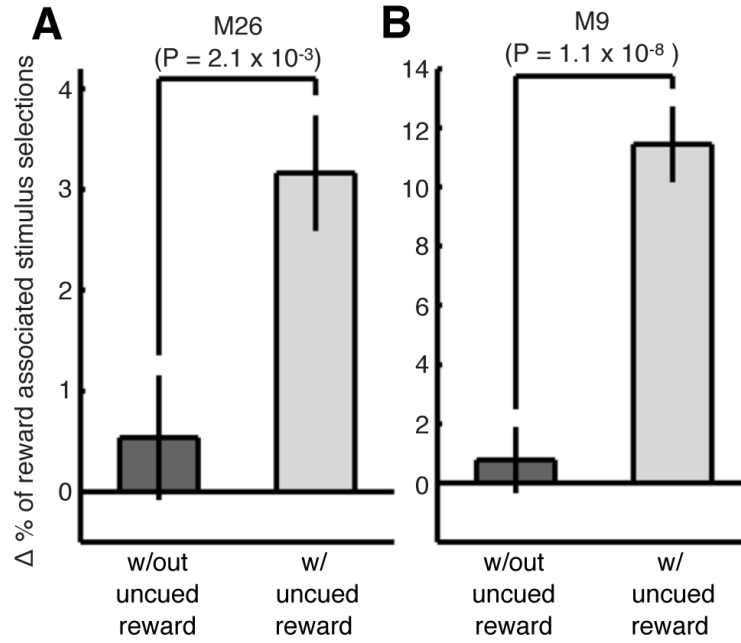


Figure 9. Uncued rewards alter the strength of cue-reward associations: behavioral evidence (experiment 7)

Changes in stimulus selection (for stimuli that were initially non-preferred for the monkey) obtained during the behavioral preference sessions. The 400 trials from the pre- and post-association preference test in each session were grouped into bins of 20 trials. The percent change in the monkey's stimulus selection for (A) M26 (26 sessions) and (B) M9 (9 sessions) was determined by comparing the number of selections within a bin *after* the association block with the number of selections within the corresponding bin *before* the association block relative to the total number of trials in the bin. This was performed for sessions with association blocks that contained uncued rewards (blue bars) and alternating sessions that did not (grey bars). Error bars denote the SEM over bins. A significant difference in stimulus selection was determined using a Wilcoxon rank sum test comparing experiments with and without uncued rewards.