The Role of Frontal Cortical and Medial-Temporal Lobe Brain Areas in Learning a Bayesian Prior Belief on Reversals

Anthony I. Jang,* Vincent D. Costa,* Peter H. Rudebeck, Yogita Chudasama, Elisabeth A. Murray, and Bruno B. Averbeck

Laboratory of Neuropsychology, National Institute of Mental Health, National Institutes of Health, Bethesda, Maryland 20892-4415

Reversal learning has been extensively studied across species as a task that indexes the ability to flexibly make and reverse deterministic stimulus–reward associations. Although various brain lesions have been found to affect performance on this task, the behavioral processes affected by these lesions have not yet been determined. This task includes at least two kinds of learning. First, subjects have to learn and reverse stimulus–reward associations in each block of trials. Second, subjects become more proficient at reversing choice preferences as they experience more reversals. We have developed a Bayesian approach to separately characterize these two learning processes. Reversal of choice behavior within each block is driven by a combination of evidence that a reversal has occurred, and a prior belief in reversals that evolves with experience across blocks. We applied the approach to behavior obtained from 89 macaques, comprising 12 lesion groups and a control group. We found that animals from all of the groups reversed more quickly as they experienced more reversals, and correspondingly they updated their prior beliefs about reversals at the same rate. However, the initial values of the priors that the various groups of animals brought to the task differed significantly, and it was these initial priors that led to the differences in behavior. Thus, by taking a Bayesian approach we find that variability in reversal-learning performance attributable to different neural systems is primarily driven by different prior beliefs about reversals that each group brings to the task.

Key words: amygdala; Bayesian prior; learning set; medial prefrontal cortex; orbitofrontal cortex; reversal learning

Significance Statement

The ability to use prior knowledge to adapt choice behavior is critical for flexible decision making. Reversal learning is often studied as a form of flexible decision making. However, prior studies have not identified which brain regions are important for the formation and use of prior beliefs to guide choice behavior. Here we develop a Bayesian approach that formally characterizes learning set as a concept, and we show that, in macaque monkeys, the amygdala and medial prefrontal cortex have a role in establishing an initial belief about the stability of the reward environment.

Introduction

The ability to flexibly alter previously learned responses is crucial for adaptive reward-guided decision making in dynamic environments. A simple and well studied task that requires flexibility in behavioral responses is object reversal learning (Mishkin, 1964; Dias et al., 1996; Murray et al., 1998; Cools et al., 2001; Roberts, 2006; Hampton et al., 2007; Rygula et al., 2010). This

task initially requires subjects to learn to choose one of two distinct objects or stimuli to gain a reward. When performance reaches an acceptable criterion, the object–reward associations are reversed and the subjects have to relearn to criterion before another reversal is instituted.

Previous studies have, however, failed to account for the fact that two separate but intertwined learning processes influence behavior during the task: (1) learning and reversing stimulus–reward associations, and (2) learning that reversals occur, known as learning set. The second of these processes occurs because, before the first reversal happens, subjects have only experienced stable object–reward associations in the context of pretraining and all previous experimental testing. Consequently, the first reversal generates unexpected uncertainty (Yu and Dayan, 2005). After reaching criterion on the first reversal, the stimulus–reward mappings are again reversed, and this process is continued. As animals experience additional reversals, they undergo a process of developing a learning set, where they learn that reversals occur

Received April 22, 2015; revised June 29, 2015; accepted July 18, 2015.

Author contributions: P.H.R., Y.C., and E.A.M. designed research; P.H.R., Y.C., and E.A.M. performed research; A.I.J., V.D.C., and B.B.A. analyzed data; A.I.J., V.D.C., P.H.R., E.A.M., and B.B.A. wrote the paper.

This work was supported by the Intramural Research Program of the National Institute of Mental Health. We thank Alicia Izquierdo, Dawn Anuszkiewicz-Lundgren, Katherine Wright, Wendy Hadfield, Robin Suda, and Anna Prescott for behavioral testing.

*A.I.J. and V.D.C. contributed equally to this work.

Correspondence should be addressed to Bruno B. Averbeck, PhD, Laboratory of Neuropsychology, NIMH/NIH, Building 49 Room 1B80, 49 Convent Drive MSC 4415, Bethesda, MD 20892-4415. E-mail: bruno.averbeck@nih.gov.
DOI:10.1523/JNEUROSCI.1594-15.2015

Copyright © 2015 the authors 0270-6474/15/3511751-10\$15.00/0

(Harlow, 1949). At this point reversals become an expected uncertainty and subjects learn to expect that reversals occur, although the exact time of the reversal is still unknown. This expectation can be formally defined as a Bayesian prior belief that a reversal in the reward contingencies will occur. The object reversal-learning task, therefore, requires the development of a prior belief about the probability that reversals occur.

Efficient performance in reversal-learning tasks has long been attributed to various regions of prefrontal and temporal cortex (Mishkin, 1964; Dias et al., 1996; Roberts, 2006). However, previous reports have not measured learning in terms of the two aforementioned processes. This means that the specific contribution of different parts of prefrontal and temporal cortex to flexible behavior is unclear. Here we directly characterize these processes using a Bayesian approach (Costa et al., 2015). We develop a model that characterizes inference of stimulus–reward associations, and show how this model is learned as the animals gain more experience with reversals. We further use this model to characterize aspects of choice behavior that change with learning of task set. We applied this approach to the reversal-learning behavior of 89 rhesus monkeys consisting of 12 lesion groups and an aggregated group of 32 unoperated controls. The large number of control animals gave us increased statistical power over typical investigations of this type. The operated groups had lesions of select regions within the medial and orbital frontal cortex, as well as of restricted regions in the medial temporal lobe (Fig. 1). All groups were tested using the same test methods and all animals were naive to reversals before the experiment, which allowed us to examine acquisition of a prior belief in reversals, with experience.

Materials and Methods

Subjects. Eighty-nine adult rhesus monkeys (*Macaca mulatta*; 85 males and 4 females) served as subjects. These animals were from 12 lesion groups (Fig. 1) and one aggregated control group. Most of the data used in our analysis were compiled from previously published selective lesion experiments (Murray et al., 1998; Izquierdo et al., 2004; Izquierdo and Murray, 2005; Saksida et al., 2007; Rudebeck and Murray, 2011; Chudasama et al., 2013; Rhodes and Murray, 2013; Rudebeck et al., 2013a, 2014). For a full description of surgical procedures and the extent of lesions, readers should consult the original papers (see above). Data from some groups are unpublished [groups area 32, area 25, and agranular orbitofrontal cortex (OFC)]. Despite their varying time frames, all experiments have been implemented under nearly identical protocols of object reversal learning, in the same laboratory, and involved monkeys not previously trained in any reversal-learning task. Thus, the training history of the groups was highly similar. All procedures were reviewed and approved by the National Institute of Mental Health Animal Care and Use Committee.

Preliminary training. Before formal training, all monkeys were habituated to the Wisconsin General Testing Apparatus by allowing them to take food *ad libitum* from the test tray. Then they were trained by successive approximation to displace objects located over the food wells to obtain food rewards. Following this preliminary training, monkeys underwent surgery that involved either excitotoxic or aspiration (ASP) lesions, or were retained as unoperated controls.

Object reversal learning. A single pair of objects, novel at the beginning of the experiment, was presented at each trial throughout object reversal

Figure 1. Intended lesion location for the 12 operated groups. Schematic illustration of the brain showing intended lesion locations.

Figure 2. Object reversal-learning task. For each trial, the monkey displaced one of two objects to gain a food reward underneath. Only one object led to a reward in a deterministic manner. When performance reached the criterion, the object–reward contingencies switched. The monkeys performed 30 trials each day, until either seven or nine serial reversals occurred.

learning (Fig. 2). Only one of the two objects led to a food reward in a deterministic manner. The objective was to learn, through trial and error, which object to displace to obtain the reward underneath. To prevent inherent object preferences from biasing initial scores, both objects were either baited (for half the monkeys in each lesion group) or unbaited (remaining monkeys) on the first trial of the acquisition phase. If the object chosen on the first trial was rewarded (i.e., when both objects were baited), it was designated as S+; if not (i.e., unbaited), it was designated as S−. Starting from trial 2, the food well underneath the object designated S+ was baited, whereas the other was not. Each trial consisted of the presentation of the pair of objects, one overlying each of the food wells of a two-well test tray. If the monkey chose the S+, it could retrieve the food reward. If the monkey chose the S−, no reward was available and the trial was terminated without correction. All monkeys performed 30 trials per day, each separated by 10 s. The left–right position of the correct object followed a pseudorandom order. The criterion was set at 93% accuracy (28 correct responses in 30 trials) on one day followed by 80% accuracy (24 correct out of 30) the next day. When the monkeys reached criterion, the object–reward contingencies reversed and the procedure was repeated, now with the previously unrewarded object being rewarded, and vice versa. The monkeys performed a total of either seven or nine serial reversals.

Data analysis. Data were organized into reversal blocks that contained 60 trials before the reversal (from the 2 d in which the previous criterion was achieved), and the subsequent trials until the next criterion had been

reached. Therefore, the blocks varied in size depending on the animal's performance, with a minimum of 120 trials (60 trials before and 60 trials after reversal). Nevertheless, all blocks were similar in that the reversal occurred on trial 61 (the first trial after the criterion had been reached). The reversals were always initiated on the day after the animals attained criterion.

The combined control group was aggregated across multiple studies. No statistically significant differences were found between the control groups from the different studies for either errors to criterion or estimated monkey reversal point ($p > 0.1$, all comparisons) and therefore the data were collapsed into one group. Reversal blocks were excluded from analysis if a monkey had accumulated >60 errors before reaching criterion. This criterion was adopted because animals either reversed with <60 errors or, in a few cases, scored far more errors before reversing. Thirteen blocks were removed as a result of this criterion [one block from controls, one block from OFC ASP lesions (OFC_{ASP}), four blocks from area 14, one block from amygdala (AMG), three blocks from unilateral OFC-AMG, three blocks from hippocampus]. Unless otherwise indicated, we always carried out analyses using data from the first seven reversals, which were common to all groups.

Bayesian analysis. We used Bayesian statistical models to characterize where the animals reversed, as well as how much evidence they had that a reversal had occurred, when they reversed. Although feedback was deterministic, the animals behaved as if they were in a stochastic environment. First, we used a model ($M = 1$, behavioral choice model) to estimate the point at which each monkey's choice preference reversed in each reversal block. This identifies when the monkey switched its behavior from mostly choosing one object to mostly choosing the other object. We also used a second model ($M = 2$, causal model) to estimate the amount of evidence the animal had accumulated regarding whether the reversal had occurred, when the monkey's choice preference reversed (i.e., at the point predicted by the first model).

The models use a similar framework that differ in details. Both models estimated the posterior probability that the reversal occurred on each trial. The likelihood function was defined as follows:

$$f(D | r, p, h, M) = \prod_{t=1}^k q(t)$$

where the variable D represents the choice and outcome data up to trial k , r is the inferred trial on which the object reversal occurred ($r \in (1, \dots, k)$), p is the animal's estimate of the probability of reward for the correct option ($p \in 0.5\text{--}0.99$, in 0.01 intervals), h is the initially rewarded object ($h \in 1, 2$), M is the model in use ($M \in 1, 2$), and t is the trial number up to the last trial of the block ($t \in (1, \dots, k)$). The analysis generates a distribution over r , and therefore an estimate of where the animal reversed its choice behavior.

The posterior probability was given by the following:

$$p(r, p, h | D, M) = \frac{f(D | r, p, h, M) p(r | M) p(p, h | M)}{p(D | M)}$$

Since the total length of each block varied depending on the monkey's performance (better performance leads to shorter blocks), the possible values of r ranged from 1 to ≥ 300 trials. For the reward probability p , we made no prior assumptions about the monkey's expected reward rate, and thus determined the posterior probability by marginalizing across all possible values of p between 0.5 and 0.99. Although p was deterministic as mentioned above, the animals behaved stochastically, and therefore internally their stimulus–reward mappings were stochastic. The priors on r , p , and h were all assumed flat. The posterior probability that the reversal occurred on trial r was calculated by marginalizing over p and h as follows:

$$p(r | D, M) = \sum_{p, h} p(r, p, h | D, M)$$

Subsequently, we determined the point estimate of the animal's reversal trial by calculating the expected value of the posterior distribution over r as follows:

$$\langle r | M = 1 \rangle = \sum_{r=0}^k p(r | D, M) r$$

For model 1, which estimated the reversal in the monkey's choice behavior, we assumed that the animals had an object preference, which they reversed at some point after the actual reversal. However, given the stable preference, they chose the wrong option in each trial with probability $1 - p$. Thus, for $t < r$ (trials before reversal) and $h = 1$ (object 1 is rewarded), choosing object 1 yields $q(t) = p$, and choosing object 2 yields $q(t) = 1 - p$ in the likelihood function. For $t \geq r$ and $h = 1$, choosing object 1: $q(t) = 1 - p$; choosing object 2: $q(t) = p$. Alternatively, if $t < r$ and $h = 2$, then choosing object 1 would yield $q(t) = 1 - p$. This was calculated using the data from the entire block because we were interested in estimating where the animals reversed.

Model 2 estimated the amount of evidence each monkey had accumulated when it switched its choice preference. In other words, the model characterized the animal's causal estimate of the probability that the reversal had occurred in the current trial, given the previous outcomes. To do so, two changes were made to model 1. First, the $q(t)$ values were determined by considering both choices and outcomes of the task and not solely the choice behavior. Therefore, when $t < r$ and $h = 1$, the possibilities were as follows: choose object 1 and get rewarded ($q(t) = p$), choose object 1 and not get rewarded ($q(t) = 1 - p$), choose object 2 and get rewarded ($q(t) = 1 - p$), and choose object 2 and not get rewarded ($q(t) = p$). For $t \geq r$ these probabilities are reversed. Second, the choice–outcome information was only available up to where the animal reversed, ($r | M = 1$). Therefore, the above values for $q(t)$ only applied to trials $t < r$. The cumulative value of the posterior up to the estimated monkey's reversal point, $p(r < k | D, M)$, was used as the posterior evidence that the reversal has occurred somewhere before the current trial.

We assumed the animals reversed their choice behavior when the relative evidence that a reversal occurred exceeded a threshold. Specifically, we calculated the odds in favor of a reversal at trial $k = \langle r | M = 1 \rangle$, where the animal reversed as follows:

$$F = \frac{p(r < k | D, B) p(B)}{(1 - p(r < k, B)) (1 - p(B))}$$

The animal's belief that reversals occur, $p(B)$, was assumed to evolve as the animals experienced more reversals. We assumed that the odds that the animals used to switch was 5. The inferred belief about reversals changed monotonically with this value, so the choice of 5 was arbitrary and the conclusions do not depend on a specific value. The evidence was calculated at $p = 0.55$ for this analysis. Again, this value is arbitrary, but the specific value does not affect our conclusions. For larger values of p , the evidence numerically saturates at 1 for several of the animals because they switched only after many unrewarded trials and this leads to problems with the numerical implementation of the analysis. Because the exact value of p only affects relative values (i.e., evidence monotonically scales with p), we chose 0.55 because this led to reasonably well behaved classical statistics. From an information point of view, however, the exact value of p does not matter for $0.5 < p < 1$. The p values for reported effects were similar up to values of $q < 0.7$.

Prior update equation. We assumed priors were updated using a delta learning rule. Specifically, we fit the following equation:

$$p(B)_t = p(B)_{t-1} + p(p(B)_\infty - p(B)_{t-1}).$$

We fit this equation by minimizing the mean squared error between the model estimate of $p(B)_t$ and the observed value for each animal. The two free parameters were $p(B)_0$, the initial value of the prior, and ρ , which was the update coefficient. The value of $p(B)_\infty$ was set to 0.95 for all animals. This could have been floated as a free parameter, but the number of data points available for fitting was only seven or nine (i.e., the number of reversals), so fixing $p(B)_\infty$ left the degrees of freedom of the model at 2, which was more robust and stable. The variable $p(B)_\infty$ is the asymptotic value of the prior.

Cluster analysis. Clustering algorithms normally proceed using various heuristics to form the clusters. For example, the K means algorithm starts with a set of K initial cluster centroids. It then assigns each item to the initial cluster centroid that it is closest to, according to some metric, and then recomputes the cluster means. This is iterated until convergence.

Table 1. Demographic information for each lesion group

Cluster: lesion group	N	Surgery	Percentage damage (range)	Number of reversals	Reference
1					
Control	32		Not applicable	7 or 9	
OFC area 14	4	EXC	79.3% (58.7–96.0%)	9	Rudebeck and Murray, 2011b
2					
OFC _{ASP} complete	3	ASP	78.5% (69.2–88.9%)	9	Izquierdo et al., 2004
Unilateral OFC _{ASP} -AMG	8	OFC: ASP AMG: EXC	OFC: 82.2% (75.3–87.4%) AMG: 98.7% (96.2–100%)	9	Izquierdo and Murray, 2004
Hippocampus	4	EXC	64% (28–83%)	9	Saksida et al. 2007 (histology; reversal data not published)
	4	EXC	51.4% (27.9–74.7%)	7	Murray et al., 1998
Rhinal cortex	3	ASP	53.0% (51.5–54.1)	7	Murray et al., 1998
3					
OFC area 11/13*	4	EXC	85.8% (81.7–91.8%)	9	Rudebeck and Murray, 2011b
Agranular OFC	3	EXC	Not estimated	9	Unpublished
AMG	5	EXC	93.9% (85.2–100%)	9	Izquierdo and Murray, 2007
Area 32* (prelimbic region of medial prefrontal cortex)	4	ASP	65.4% (55.0–82.6%)	7	Rhodes and Murray, 2013 (histology; reversal data not published)
Area 25* (subgenual region of medial prefrontal cortex)	4	ASP	Not reported	7	Rudebeck et al., 2014 (histology; reversal data not published)
Area 24* (anterior cingulate portion of medial prefrontal cortex)	4	ASP	85.5% (72.0–96.2%)	7	Chudasama et al., 2012
OFC _{EXC} complete	7	EXC	77.2% (64.3–96.3%)	9	Rudebeck et al. 2013

EXC, excitotoxic.

*Walker's area 11/13, 14 and Carmichael and Price areas 24, 25, 32.

There is no principled way to choose the initial cluster centroids, and the initial centroids affect the final group membership. This can be overcome by exhaustively assigning each group to all of the different clusters, computing centroids for each assignment, and then using the clustering that led to the smallest distances to the centroids. This approach supersedes all clustering algorithms as it is guaranteed to identify the “best” clustering, and it is not subject to local minima. Although this approach is generally computationally prohibitive, it is often possible to examine a substantial fraction of the clusters to see which of them lead to small distances to the centroids (Averbeck and Seo, 2008). This is the approach we took in the current study. Further, we used the *F* value from an ANOVA analysis as our distance metric. This penalizes degrees of freedom (i.e., number of clusters) in a principled way.

In detail, we proceeded as follows. We randomly assigned the 13 groups to between two and four clusters. We then carried out a mixed-effects ANOVA with cluster membership as a fixed factor. Lesion group was nested under cluster as a fixed effect, and monkey was nested under lesion group as a random effect. Reversal trial was also a fixed effect and the dependent variable was the estimated reversal point. We then computed the sum of the *F* value for the main effect of cluster and an interaction between cluster and reversal, and used this as our distance metric. This summarizes all effects of cluster. We examined 9,000,000 random clusterings out of a total of 4^{13} possible clusterings (i.e., each of the 13 groups gets randomly assigned to one cluster, which is a 13-digit quaternary string = 4^{13}). This is a reasonable fraction of the maximum number of possible clusters and therefore should be useful for finding the robust clusters in the data.

We next determined which clusters occurred most frequently in the best random cluster groups. This is similar to computing an expected value in cluster space. To find the clusters that occurred most frequently, we sorted the clusterings by the combined *F* statistic from the ANOVA and picked the top 1000 models, where a model is defined by its clustering (i.e., which cluster each of the groups was assigned to). We then found the individual clusters that occurred most frequently in the top 1000 clusters (for example, how often do Groups 1, 5, and 8 cluster together). Finally, we found sets of clusters that grouped all of the lesions (i.e., by exhaustively combining the clusters that occurred in the top 1000 to form sets of clusters that included exactly all of the 13 groups once), and found which of these sets of clusters had the highest frequency, from among the sets that occurred in the top 1000 clusters. Here, we report the

clusters that occurred most frequently, using this analysis. We also carried out the same analysis by finding the cluster sets that had the highest total summed *F* value. However, this gave the same answer because there was less variation in the *F* value across the top 1000 lesion groups than there was in the frequency of occurrence of clusters. Note that the clustering that we report in the Results had the fourth highest *F* in our sample, overall.

Classical statistical analyses. Behavioral performance was analyzed using mixed-effects ANOVA models. Monkey was modeled as a random effect nested under lesion group. In the analyses that used cluster as a factor, lesion group was further nested under cluster, with the individual monkeys nested under lesion. Reversal number was always modeled as a continuous covariate. For the analysis of reversal in individual groups, we used the maximum number of reversals available (seven or nine reversals). Generally the improvement in performance was reasonably well approximated as linear.

Results

Bayesian analysis of reversal learning

Monkeys from 12 lesion groups and an unoperated control group (Table 1; Fig. 1) carried out an object reversal-learning task (Fig. 2). Several operated groups had lesions of OFC, including groups with bilateral ASP lesions of OFC (11, 13, and 14), bilateral excitotoxic lesions of OFC (11, 13, and 14), components of OFC including areas 11/13 and area 14 (separately), and caudal agranular OFC. Within medial frontal cortex, separate groups of monkeys sustained ASP lesions of areas 24, 25, and 32. Within the medial temporal lobe, separate groups of monkeys sustained lesions of the rhinal (i.e., entorhinal and perirhinal) cortex, the hippocampus, and the AMG (Fig. 1). A group with combined unilateral ASP lesions of OFC and excitotoxic lesions of the AMG was also included.

We analyzed reversal behavior using a Bayesian model that inferred where the animals switched their choice behavior, following a reversal in the stimulus–reward contingencies. The algorithm generates a posterior distribution over the trial on which the switch in choice behavior occurred (Fig. 3A). The expected value of this distribution gives a point estimate of the

Figure 3. Choice behavior of the control group determined by the behavioral choice model ($M = 1$). **A**, The posterior probability distribution of all nine reversals. **B**, The estimated reversal point determined by the behavioral choice model and total errors to criterion for each reversal. The estimated reversal point is subtracted by 61 to allow the zero point to indicate the point of actual reversal.

Figure 4. Estimated reversal points of all lesion groups and unoperated controls, classified into three clusters. **A**, Excitotoxic lesion of medial OFC (Walker's area 14) clustered together with the control group, showing no signs of deficit. **B**, The OFC_{ASP}, rhinal cortex, hippocampus, and combined unilateral AMG-OFC lesion groups clustered together, showing varying degrees of impairment in detecting reversals. **C**, The AMG and cingulate region groups (areas 24, 25, 32), as well as groups with excitotoxic lesions of lateral OFC (Walker's area 11/13) and agranular insular cortex showed a tendency of reversing earlier compared with controls in the first few reversals.

trial on which the animal switched its choice behavior (Fig. 3B). The reversal point often provides a measure consistent with the errors to criterion (Fig. 3B), which has been used previously to analyze reversal behavior. However, analyzing the data using the Bayesian model allowed us to characterize the improvement in reversal performance as the animals experienced more reversals as an increase in prior beliefs that reversals occur in the world, relative to the evidence accumulated.

We first considered an analysis of the errors to criterion across our lesion groups. We found that there was a main effect of lesion ($F_{(12,76)} = 3.94$, $p < 0.001$), a main effect of reversal ($F_{(1,76)} = 71.14$, $p < 0.001$) but no lesion-by-reversal interaction ($F_{(12,76)} = 1.09$, $p = 0.384$). Thus, there were overall differences in trials to criterion, and these improved with experience, but there were no group differences in the rates at which they improved across groups, when estimated with the trials to criterion.

We next compared reversal performance in the 13 groups using the Bayesian model to identify the estimated reversal point and used this as the dependent variable. Across the 13 groups (Fig. 4; plotted in clusters of groups, which we define below) there were main effects of lesion ($F_{(12,77)} = 4.85$; $p < 0.001$) and reversal ($F_{(1,77)} = 114.81$; $p < 0.001$), indicating that groups differed in how quickly they reversed their choice behavior, and that they generally reversed more quickly as they

had more experience with reversals. There was also a lesion-by-reversal interaction ($F_{(12,77)} = 2.94$; $p = 0.002$), indicating that groups varied in the rate at which their performance changed across reversals.

Next, we carried out analyses in individual lesion groups to see whether the estimated reversal trial decreased with reversals. This would indicate that the animals were learning to reverse on the basis of less evidence, i.e., developing a reversal-learning set. These analyses were carried out on the total number of reversals administered to each group (seven or nine reversals). We found that reversal was significant in each individual group (all $p \leq 0.028$), except the OFC_{ASP} ($p = 0.132$) group. One of the animals in the OFC_{ASP} group had poor performance on intermediate reversals. Therefore, we compared the data from the first reversal with the last in the OFC_{ASP} group (to see whether there was improvement from beginning to end) and found that there was a significant difference (Mann–Whitney U , $p = 0.05$). In summary, there was evidence that all of the groups improved their performance with reversal experience.

Cluster analysis

Upon inspection of the estimated reversal point performance across groups, it appeared that performance clustered such that certain lesions led to similar performance. Therefore, we carried out a cluster analysis to determine which groups clustered to-

Figure 5. Estimated reversal points and posterior probability distributions averaged across lesion groups for each cluster. **A**, Cluster 2 showed a strong overall impairment, while Cluster 3 showed an enhancement in learning to reverse, especially in the first few reversals. Regardless, all clusters still learned to reverse earlier with more experience. **B**, The posterior distribution shows a marked enhancement of cluster 3 in both accuracy and precision during the first reversal.

gether (see Materials and Methods). The analysis also allows for a more compact summary of the data. We found that the data supported three clusters (Fig. 4), despite the fact that we allowed for ≤ 4 clusters in our analysis. Specifically, one of the clusters was composed of the lesion groups with excitotoxic lesions of medial OFC (Walker's area 14) and controls (cluster 1; Fig. 4A). From here on, we refer to these groups as cluster 1. Another cluster contained all the lesion groups that showed a deficit compared with controls, which included the OFC_{ASP}, rhinal cortex, hippocampus, and combined unilateral AMG-OFC lesion groups (cluster 2; Fig. 4B). From here on, we refer to these lesion groups as cluster 2. The final cluster included the group with AMG lesions, groups with excitotoxic lesions of the entire OFC, the lateral OFC (Walker's area 11/13), and agranular OFC, and the three groups with lesions that fell within the medial frontal cortex (areas 24, 25, 32; cluster 3; Fig. 4C). Lesion groups in this cluster reversed faster than controls in the first few reversals, and thus overall their performance was better on the task. From here on, we refer to these lesion groups as cluster 3.

We next characterized the statistics of the clusters. We analyzed the estimated reversal trial using cluster as a factor with lesion nested under cluster (Fig. 5). There were significant effects of cluster ($F_{(2,77)} = 24.59$; $p < 0.001$), reversal ($F_{(1,77)} = 142.44$; $p < 0.001$), and a cluster-by-reversal interaction ($F_{(2,76)} = 11.53$; $p < 0.001$). Furthermore, all pairwise comparisons among the three clusters revealed a significant effect of cluster ($p < 0.01$, all comparisons, uncorrected). In addition, cluster 3 reversed more quickly over the first few reversals relative to cluster 1, and the cluster-by-reversal interaction was significant when cluster 3 was compared with cluster 1 ($F_{(1,58)} = 9.30$, $p = 0.004$, uncorrected). Finally, when each cluster was analyzed individually, they all improved over reversals ($p < 0.001$). Thus, even cluster 2 improved with reversal, when group OFC_{ASP} was combined with the other impaired lesion groups. When the posterior distribution over reversals was examined, it could be seen that cluster 3 reversed much more effectively relative to the others on the first reversal (Fig. 5B). These effects are not surprising, given the individual lesion group effects reported above, and the separation of clusters on the basis of their reversal behavior.

Analysis of choices

To evaluate potential differences in choice consistency following reversals, we examined each cluster's choice behavior aligned at the actual reversal point and compared this with alignment at the monkey's estimated reversal point (Fig. 6). To correct for differences in the number of trials before and after each reversal, we split each block into two phases (i.e., prereversal and postreversal phases). Then we normalized the number of trials in both phases by creating 10 equally sized bins with intervals appropriately scaled to the number of trials in each phase. This allowed us to plot each group's choice behavior aligned at the estimated reversal point. When the curves were aligned at the actual reversal point, the average choice behavior shifts slowly across trials, and it appears that there are several trials during which the animals are stochastically sampling both options, particularly in the first reversal (Fig. 6, dashed red lines). This apparent gradual shift in the average choice behavior is also consistent, however, with each animal switching quickly, but at different delays relative to the reversal trial. To disambiguate this, we also examined average choice behavior with the data aligned to the estimated reversal trial for each animal. This analysis aligns switches that occur at different delays across animals. When the data are examined this way, the switches were more abrupt, and it can be seen that once the animals switch, they rarely sample the other option in subsequent trials. This was generally true for both the first and the final reversal (Fig. 6, dashed and solid blue lines). Thus, even during the first reversal, the animals sampled minimally among the two choices while reversing.

To further investigate variations in the monkeys' strategies, we characterized their choices by examining how often they chose the same option that had just been rewarded (win–stay) and how often they switched options after not receiving a reward (lose–shift). Because the feedback was deterministic, a win–stay/lose–shift strategy would be optimal in this task, although the animals were not behaving optimally. This allowed us to dissociate differences in the pattern of errors committed by each group and to determine the extent to which choice was influenced by the outcome of the previous trial. For each monkey, we calculated the conditional probabilities of win–stay [receiving a reward followed by choosing the same option, $p(\text{stay} | \text{win})$], and lose–shift

Figure 6. Choice behavior aligned at the estimated reversal point shows that monkeys exhibit minimal sampling behavior. Although aligning the choice behavior at the actual reversal point (red) may suggest a reinforcement learning process after the switch, aligning at the estimated reversal point (blue) shows that monkeys rarely sample the other option once they suspect a reversal has occurred. Trial numbers are normalized by creating 10 bins before and after the aligned reversal point.

Figure 7. Win-stay/lose-shift behavior of the three clusters. **A**, Cluster 2 displayed less win-stay behavior compared with the other two clusters, but the probability of win-stay remained high and stable across all reversals for all three clusters. **B**, Lose-shift behavior differed between clusters, but this behavior also increased as the monkeys experienced more reversals.

(receiving no reward followed by choosing the other option, $p(\text{shift} | \text{lose})$) for each block, using the data after the reward contingencies were reversed (Fig. 7). Clusters 1 and 3 had the most consistent win-stay behavior, followed by cluster 2. There was a significant effect of cluster ($F_{(2,77)} = 7.42$; $p = 0.0011$) but not reversal ($F_{(1,77)} = 2.21$, $p = 0.141$). Thus, the groups did not tend to improve their win-stay performance. Lose-shift performance across clusters also differed ($F_{(2,79)} = 4.56$; $p = 0.013$). In this case there was an effect of reversal ($F_{(1,77)} = 163.80$; $p < 0.001$), but no interaction ($F_{(2,77)} = 0.53$, $p = 0.59$). Thus, the clusters differed in both win-stay and lose-shift behavior, but only lose-shift behavior improved as the animals experienced more reversals. In addition, the win-stay probabilities were high in most groups. This further shows that once animals sampled the previously unrewarded option and received a reward, they tended to stick with it.

Evidence and prior

As monkeys experienced additional reversals, they reversed more quickly. We sought to characterize this phenomenon—reversal-learning set—in more detail. The Bayesian model is a specific hypothesis about how the monkeys solve this task. The model

assumes that when the animals pick one of the options and do not get rewarded, this is evidence that they should shift their current object choice. However, if they have a prior belief that stable choice strategies are more appropriate, then they will persist longer before shifting, because the evidence has to overcome this prior belief. To characterize this, we examined how prior beliefs about reversals evolved as the animals gained more experience with the task.

The monkeys reversed their object choices after experiencing a series of unrewarded choices. These unrewarded choices were evidence that their current object choice was no longer correct. The number of unrewarded choices that the animals experienced before they switched their object choice, and therefore the evidence that a reversal had occurred, decreased as they had more experience with the reversals (Fig. 8A; main effect of reversal $F_{(1,77)} = 220.44$, $p < 0.001$) and this effect differed across clusters ($F_{(2,77)} = 23.87$, $p < 0.001$). Therefore, although the three clusters differed in the overall evidence required before deciding to reverse, all clusters learned to reverse with less evidence as they experienced more reversals. The ANOVA also revealed a cluster-by-reversal interaction ($F_{(2,77)} = 7.59$; $p < 0.001$).

Figure 8. Evidence and prior at each reversal point determined by the causal model ($M = 2$). **A**, Evidence accumulated regarding whether the reversal had occurred when the monkey decided to reverse its behavior. All clusters learned to reverse with less evidence with more experience. **B**, Prior belief of the monkey that reversals occur in the environment. As they experience more reversals, all clusters increased their prior belief that reversals occur. Solid lines are data; dashed lines are model fits to data.

The monkey's decision to reverse its choice preference depended on both the evidence that the specific object signaling reward availability had changed (Fig. 8A) and prior beliefs that stimulus–reward mappings reverse (Fig. 8B). The evidence in favor of a reversal can be summarized by calculating the odds in favor of a reversal (i.e., either a reversal has or has not occurred). We assumed that animals reversed when the odds favoring a reversal exceeded a fixed threshold. Using a fixed threshold, we calculated their belief that reversals occur during each reversal block (Fig. 8B). The prior varied by cluster ($F_{(2,79)} = 15.81$; $p < 0.001$) and reversal ($F_{(1,78)} = 215.16$; $p < 0.001$), although there was no cluster-by-reversal interaction ($F_{(2,77)} = 2.10$; $p = 0.130$). In sum, whereas the three clusters differed in their overall belief that reversals occur, they all developed a stronger belief that reversals occur as they gained experience. We also quantified the initial value of the prior, and the prior update with experience, using a learning model (Fig. 8B). We found that the initial values varied among clusters ($F_{(2,86)} = 7.46$, $p = 0.001$) but that the rate at which priors were updated, as the animals experience more reversals, did not vary among clusters ($F_{(2,86)} = 2.28$, $p = 0.109$). *Post hoc* comparisons showed that the initial value varied significantly between cluster 2 and cluster 3 ($p = 0.003$, uncorrected) and marginally between clusters 1 and 3 ($p = 0.032$, uncorrected) but not between clusters 1 and 2 ($p = 0.056$, uncorrected). Thus, the clusters differed in how they initially approached the task, but they all learned about reversals at the same rate.

In a final analysis we examined the priors among the individual lesion groups. We found an overall effect of group on the prior ($F_{(12,76)} = 2.9$, $p = 0.002$). We followed this up with the 72 pairwise *post hoc* comparisons among groups, Bonferroni corrected. Significance, therefore, required a p value of $0.05/78 = 6.4 \times 10^{-4}$. We found only three significant pairwise comparisons at this level. Specifically, the rhinal group differed from the area 32 group, the area 25 group, and the area 24 group.

Discussion

To determine how different parts of the frontal cortex and medial temporal lobe contribute to flexible reward-guided behavior as measured by object reversal learning, this study examined the choice behavior of 89 rhesus monkeys comprising 13 groups. This manuscript reports on data from several lesion groups that have not been previously published (Table 1), and also develops a

novel Bayesian approach, which clarifies the changes in behavior that underlie the development of learning set, while also providing a formal interpretation.

Bayesian characterization of reversal behavior

We used a Bayesian-modeling approach to differentiate hypotheses that could account for previously published results that used errors to criterion as a performance measure. Specifically, errors to criterion can improve within a group for one of two reasons. First, a group may consistently pick one stimulus and, following the reversal, continue to pick that stimulus for several trials before switching to the other stimulus and then consistently picking that stimulus. Their performance would improve in this case if they perseverated on the unrewarded stimulus for fewer trials as they gained experience with reversals. Alternatively, following the reversal the groups may have stochastically chosen both options for several trials, before choosing the currently rewarded stimulus more consistently. Simply examining average learning curves or errors to criterion cannot differentiate these possibilities. We show that the first strategy characterizes the behavior more accurately. Specifically, when we use the Bayesian model to identify the reversal point, and average choice behavior around that point, we find that the animals were switching abruptly, as opposed to stochastically sampling both options for a period. Second, when we estimated the win–stay/lose–shift performance, we found that learning set corresponded with improvement in lose–shift, but not win–stay. Win–stay performance was generally high for all groups, and did not improve with experience. Thus, the acquisition of learning set was due specifically to a reduction in the number of perseverative trials before animals switched preferences, and not an improvement in the variability of choice behavior following a reversal.

The Bayesian model also provides a formal interpretation of learning set, as acquisition of a prior on reversals (Harlow, 1949). As the monkeys gained experience with the task they switched after fewer unrewarded trials. These unrewarded trials are evidence that the stimulus–reward mapping has switched. This evidence combined with a prior belief about whether reversals occur generates a posterior belief about whether a reversal has occurred. When this posterior belief exceeds a threshold, the animals switch their object choice. Assuming the threshold is constant, these changes in prior beliefs reflect the acquisition of a learning set.

When we separately characterized the initial value of the prior beliefs, and the rate at which these were updated across groups, we found that the initial values varied across groups, but the rate at which they were updated did not. Thus, none of the lesions affected the rate of update of the prior belief, and correspondingly acquisition of learning set. However, they did affect the initial values of the prior beliefs.

Differential effects of brain lesions on reversal behavior

As has been appreciated from previous studies, ASP lesions of OFC (Butter, 1969; Iversen and Mishkin, 1970; Meunier et al., 1997; Izquierdo et al., 2004), as well as lesions of rhinal cortex, excitotoxic lesions of the hippocampus, and unilateral lesions of OFC and the amygdala within one hemisphere (Murray et al., 1998; Izquierdo et al., 2004) all produce significant impairments in reversal learning. Prior studies testing for visual perceptual deficits in monkeys with rhinal cortex lesions (Gaffan and Murray, 1992; Meunier et al., 1993) suggest that visual perceptual abilities are not the primary cause of the impairments in reversal learning for monkeys in this lesion cluster. Here, we also show these deficits are not due to an inability of the animals to develop a learning set and are instead attributable to a need for greater evidence in favor of a reversal before the animals switched their behavior.

In contrast to the cluster with deficits relative to controls, we also identified a cluster of lesion groups that performed better than controls on the first few reversals. This cluster included a group with excitotoxic lesions of the amygdala together with several groups with lesions in prefrontal cortical areas that receive substantial, direct inputs from the AMG (Ghashghaei et al., 2007) and entorhinal cortex (Saleem et al., 2008). The performance of this cluster raises the question of why they reverse more quickly for the first few reversals. There are at least two possibilities, which may not apply equally to all lesion groups. First, they may have a stronger prior on reversals, consistent with what we have shown using the Bayesian model. This may also more generally be a stronger prior that the world is unstable, and that stimulus–reward mappings do not tend to be consistent over time. In accord with this idea, Murray et al. (2011) proposed that agranular frontal cortex areas (such as areas 25, 32, 24, and agranular OFC) may provide a form of top-down control that serves to bias choices toward newer memories. In the present case, intact agranular frontal areas would bias monkeys toward the most recently learned object–reward association. In the absence of this bias, monkeys would be expected to reverse more quickly.

Second, they may form weaker associations between stimuli and rewards. Inconsistent with this is the fact that this group has the strongest win–stay performance. Consistent with this, however, is the finding that monkeys with AMG lesions extinguish responding more quickly than controls in instrumental (Izquierdo and Murray, 2005; Clarke et al., 2008; Rhodes and Murray, 2013) and conditioned reinforcement (Parkinson et al., 2001) paradigms. In several settings, AMG damage has been shown to disrupt Pavlovian stimulus–outcome associations (Everitt et al., 2003; Balleine and Killcross, 2006). At the very least, AMG lesions attenuate anticipated reward coding in other brain areas at the time of choice (Rudebeck et al., 2013b). Thus, the faster reversal learning during the early reversals may reflect a reduction in anticipatory reward coding for the previously rewarded object, which would normally be expected to support choice of that object. It is also possible that a weakened ability to form stimulus–reward associations leads to stronger priors on instability. If the animal is unable to learn consistent statistical relationships in the environment, it may assume that the environment is unstable.

There are few studies of the effects of lesions on object reversal learning in medial frontal areas, and the finding of facilitatory effects of frontal cortex lesions on deterministic object reversal learning has not been previously reported. The AMG projects to the three medial frontal areas we studied (Ghashghaei et al., 2007), so one possible explanation for this finding is that the facilitation of object reversal learning is related to the AMG's influences on these cortical regions. Against this idea, however, is the fact that the AMG also projects to the inferior frontal convexity (Porrino et al., 1981), a region that, when damaged, produces impairments in object reversal (Iversen and Mishkin, 1970). An alternative possibility is that damage to medial frontal areas, primarily area 24, reduces the influence of action–reward strategies on monkeys' behavior, augmenting stimulus–reward processes. This would fit with reports that lesions of medial frontal cortex produce deficits in action-based, but not stimulus-based, learning (Rudebeck et al., 2008). Neither of these possibilities can fully explain our data and divergence from previous reports suggesting a slight decrement in performance after lesions of area 24 (Chudasama et al., 2013). This means that future studies are required to determine the role of medial frontal areas in flexible stimulus–reward learning.

Conclusion

Using a novel Bayesian approach, we compared the adaptive choice of unoperated controls and monkeys with different lesions within the frontal cortex and temporal lobe. First, we found that across the population of monkeys, reversal behavior clustered into three groups. ASP lesions involving orbital frontal and rhinal cortex, as well as excitotoxic lesions of the hippocampus, led to deficits on the task. Excitotoxic lesions of medial OFC area 14 had no effect on task performance, whereas lesions to the AMG, as well as the areas in the OFC that receive a strong AMG input and parts of the medial prefrontal cortex, led to improved performance on the first few reversals. Second, we found that reversal behavior was consistent with abrupt switching across all groups, as opposed to a period of sampling both options stochastically. Thus, the animals adopted a strategy of primarily sticking with one option, and then switching to the other option. Third, when we quantified this effect with a Bayesian model, we found that the clusters differed in the initial value of their prior beliefs in reversals, but that all clusters updated their belief in reversals at the same rate.

References

- Averbeck BB, Seo M (2008) The statistical neuroanatomy of frontal networks in the macaque. *PLoS Comput Biol* 4:e1000050. [CrossRef Medline](#)
- Balleine BW, Killcross S (2006) Parallel incentive processing: an integrated view of amygdala function. *Trends Neurosci* 29:272–279. [CrossRef Medline](#)
- Butter CM (1969) Perseveration in extinction and in discrimination reversal tasks following selective frontal ablations in *Macaca mulatta*. *Physiol Behav* 4:163–171. [CrossRef](#)
- Chudasama Y, Daniels TE, Gorrin DP, Rhodes SE, Rudebeck PH, Murray EA (2013) The role of the anterior cingulate cortex in choices based on reward value and reward contingency. *Cereb Cortex* 23:2884–2898. [CrossRef Medline](#)
- Clarke HF, Robbins TW, Roberts AC (2008) Lesions of the medial striatum in monkeys produce perseverative impairments during reversal learning similar to those produced by lesions of the orbitofrontal cortex. *J Neurosci* 28:10972–10982. [CrossRef Medline](#)
- Cools R, Barker RA, Sahakian BJ, Robbins TW (2001) Enhanced or impaired cognitive function in Parkinson's disease as a function of dopaminergic medication and task demands. *Cereb Cortex* 11:1136–1143. [CrossRef Medline](#)
- Costa VD, Tran VL, Turchi J, Averbeck BB (2015) Reversal learning and

- dopamine: a Bayesian perspective. *J Neurosci* 35:2407–2416. [CrossRef Medline](#)
- Dias R, Robbins TW, Roberts AC (1996) Dissociation in prefrontal cortex of affective and attentional shifts. *Nature* 380:69–72. [CrossRef Medline](#)
- Everitt BJ, Cardinal RN, Parkinson JA, Robbins TW (2003) Appetitive behavior: impact of amygdala-dependent mechanisms of emotional learning. *Ann N Y Acad Sci* 985:233–250. [Medline](#)
- Gaffan D, Murray EA (1992) Monkeys (*Macaca fascicularis*) with rhinal cortex ablations succeed in object discrimination learning despite 24-h intertrial intervals and fail at matching to sample despite double sample presentations. *Behav Neurosci* 106:30–38. [CrossRef Medline](#)
- Ghashghaei HT, Hilgetag CC, Barbas H (2007) Sequence of information processing for emotions based on the anatomic dialogue between prefrontal cortex and amygdala. *Neuroimage* 34:905–923. [CrossRef Medline](#)
- Hampton AN, Adolphs R, Tyszka MJ, O'Doherty JP (2007) Contributions of the amygdala to reward expectancy and choice signals in human prefrontal cortex. *Neuron* 55:545–555. [CrossRef Medline](#)
- Harlow HF (1949) The formation of learning sets. *Psychol Rev* 56:51–65. [CrossRef Medline](#)
- Iversen SD, Mishkin M (1970) Perseverative interference in monkeys following selective lesions of the inferior prefrontal convexity. *Exp Brain Res* 11:376–386. [Medline](#)
- Izquierdo A, Murray EA (2005) Opposing effects of amygdala and orbital prefrontal cortex lesions on the extinction of instrumental responding in macaque monkeys. *Eur J Neurosci* 22:2341–2346. [CrossRef Medline](#)
- Izquierdo A, Suda RK, Murray EA (2004) Bilateral orbital prefrontal cortex lesions in rhesus monkeys disrupt choices guided by both reward value and reward contingency. *J Neurosci* 24:7540–7548. [CrossRef Medline](#)
- Meunier M, Bachevalier J, Mishkin M, Murray EA (1993) Effects on visual recognition of combined and separate ablations of the entorhinal and perirhinal cortex in rhesus monkeys. *J Neurosci* 13:5418–5432. [Medline](#)
- Meunier M, Bachevalier J, Mishkin M (1997) Effects of orbital frontal and anterior cingulate lesions on object and spatial memory in rhesus monkeys. *Neuropsychologia* 35:999–1015. [CrossRef Medline](#)
- Mishkin M (1964) Perseveration of central sets after frontal lesions in monkeys. In: *The frontal granular cortex and behavior* (Warren JM, Akert K, eds), pp 219–241. New York: McGraw-Hill.
- Murray EA, Baxter MG, Gaffan D (1998) Monkeys with rhinal cortex damage or neurotoxic hippocampal lesions are impaired on spatial scene learning and object reversals. *Behav Neurosci* 112:1291–1303. [CrossRef Medline](#)
- Murray EA, Wise SP, Rhodes SEV (2011) What can different brains do with reward? In: *Neurobiology of sensation and reward* (Gottfried JA, ed), pp 57–92. Boca Raton, FL: CRC.
- Parkinson JA, Crofts HS, McGuigan M, Tomic DL, Everitt BJ, Roberts AC (2001) The role of the primate amygdala in conditioned reinforcement. *J Neurosci* 21:7770–7780. [Medline](#)
- Porrino LJ, Crane AM, Goldman-Rakic PS (1981) Direct and indirect pathways from the amygdala to the frontal lobe in rhesus monkeys. *J Comp Neurol* 198:121–136. [CrossRef Medline](#)
- Rhodes SE, Murray EA (2013) Differential effects of amygdala, orbital prefrontal cortex, and prelimbic cortex lesions on goal-directed behavior in rhesus macaques. *J Neurosci* 33:3380–3389. [CrossRef Medline](#)
- Roberts AC (2006) Primate orbitofrontal cortex and adaptive behaviour. *Trends Cogn Sci* 10:83–90. [CrossRef Medline](#)
- Rudebeck PH, Murray EA (2011) Dissociable effects of subtotal lesions within the macaque orbital prefrontal cortex on reward-guided behavior. *J Neurosci* 31:10569–10578. [CrossRef Medline](#)
- Rudebeck PH, Behrens TE, Kennerley SW, Baxter MG, Buckley MJ, Walton ME, Rushworth MF (2008) Frontal cortex subregions play distinct roles in choices between actions and stimuli. *J Neurosci* 28:13775–13785. [CrossRef Medline](#)
- Rudebeck PH, Saunders RC, Prescott AT, Chau LS, Murray EA (2013a) Prefrontal mechanisms of behavioral flexibility, emotion regulation and value updating. *Nat Neurosci* 16:1140–1145. [CrossRef Medline](#)
- Rudebeck PH, Mitz AR, Chacko RV, Murray EA (2013b) Effects of amygdala lesions on reward-value coding in orbital and medial prefrontal cortex. *Neuron* 80:1519–1531. [CrossRef Medline](#)
- Rudebeck PH, Putnam PT, Daniels TE, Yang T, Mitz AR, Rhodes SE, Murray EA (2014) A role for primate subgenual cingulate cortex in sustaining autonomic arousal. *Proc Natl Acad Sci U S A* 111:5391–5396. [CrossRef Medline](#)
- Rygula R, Walker SC, Clarke HF, Robbins TW, Roberts AC (2010) Differential contributions of the primate ventrolateral prefrontal and orbitofrontal cortex to serial reversal learning. *J Neurosci* 30:14552–14559. [CrossRef Medline](#)
- Saksida LM, Bussey TJ, Buckmaster CA, Murray EA (2007) Impairment and facilitation of transverse patterning after lesions of the perirhinal cortex and hippocampus, respectively. *Cereb Cortex* 17:108–115. [Medline](#)
- Saleem KS, Kondo H, Price JL (2008) Complementary circuits connecting the orbital and medial prefrontal networks with the temporal, insular, and opercular cortex in the macaque monkey. *J Comp Neurol* 506:659–693. [CrossRef Medline](#)
- Yu AJ, Dayan P (2005) Uncertainty, neuromodulation, and attention. *Neuron* 46:681–692. [CrossRef Medline](#)