Decoding the neural substrates of reward-related decision making with functional MRI

Alan N. Hampton, and John P. O'Doherty

This information is current as of January 2007.

Supplementary Material	Supplementary material can be found at: www.pnas.org/cgi/content/full/0606297104/DC1
	This article has been cited by other articles: www.pnas.org#otherarticles
E-mail Alerts	Receive free email alerts when new articles cite this article - sign up in the box at the top right corner of the article or click here.
Rights & Permissions	To reproduce this article in part (figures, tables) or in entirety, see: www.pnas.org/misc/rightperm.shtml
Reprints	To order reprints, see: www.pnas.org/misc/reprints.shtml

Notes:

Decoding the neural substrates of reward-related decision making with functional MRI

Alan N. Hampton[†] and John P. O'Doherty[†‡§]

[†]Computation and Neural Systems Program, [‡]Division of Humanities and Social Sciences, California Institute of Technology, 1200 East California Boulevard, M/C 228-77, Pasadena, CA 91125

Although previous studies have implicated a diverse set of brain regions in reward-related decision making, it is not yet known which of these regions contain information that directly reflects a decision. Here, we measured brain activity using functional MRI in a group of subjects while they performed a simple reward-based decision-making task: probabilistic reversal-learning. We recorded brain activity from nine distinct regions of interest previously implicated in decision making and separated out local spatially distributed signals in each region from global differences in signal. Using a multivariate analysis approach, we determined the extent to which global and local signals could be used to decode subjects' subsequent behavioral choice, based on their brain activity on the preceding trial. We found that subjects' decisions could be decoded to a high level of accuracy on the basis of both local and global signals even before they were required to make a choice, and even before they knew which physical action would be required. Furthermore, the combined signals from three specific brain areas (anterior cingulate cortex, medial prefrontal cortex, and ventral striatum) were found to provide all of the information sufficient to decode subjects' decisions out of all of the regions we studied. These findings implicate a specific network of regions in encoding information relevant to subsequent behavioral choice.

anterior cingulate cortex | distributed encoding | prefrontal cortex

Decision making is a neural process that intervenes between the processing of a stimulus input and the generation of an appropriate motor output. Motor responses are often performed to obtain reward, and the obligation of a decision-making mechanism is to ensure that appropriate responses are selected to maximize available reward. Although the neural systems involved in this process have been the subject of much recent research, studies have yet to isolate the specific neural circuits responsible for this decision process. Neural signals have been found that relate to but do not directly reflect this process, such as those pertaining to the expected value or utility of the available actions (1, 2), responses signaling errors in those predictions (3), encoding the value of outcomes received (4), and responses related to monitoring or evaluation of a previously executed action (5–7). Such signals have been found in diverse regions throughout the brain, including anterior cingulate cortex (ACC), medial prefrontal cortex (mPFC), orbitofrontal cortex, dorsolateral prefrontal cortex (DLPFC), amygdala, and striatum. Although complex behavioral decisions are likely to depend on information computed in a widely distributed network, it is not yet known where among this network of brain regions neural activity directly reflects the subsequent behavioral decision as to which action to select.

To determine the brain regions where neural activity is directly related to a final behavioral decision, we applied multivariate decoding techniques to our functional MRI (fMRI) data. This approach combines the temporal and spatial resolution of event-related fMRI with statistical learning techniques to decode on a trial by trial basis subjects' behavior or subjective states directly from their neural activity. Up to now, this technique has been used in visual perception, to decode perceptual states and/or perceptual decisions from fMRI signals recorded mainly (although not exclusively) in visual cortical areas (8–12). These previous studies have used locally distributed variations in activity to decode visual percepts, under situations where the global mean signals in a given region may show no significant differences between conditions. In the present case, many of our target regions of interest have been found to show global signal changes related to behavioral choice, that is, large spatially extended cluster areas of activation have previously been reported in these areas in previous fMRI studies (13, 14). Here, in addition to testing for global signals, we also tested for the presence of locally distributed signals relevant to behavioral decision making in each of our areas of interest. For this, we separated out global and local signals within each region and explored the separate contributions of signals at these two different spatial scales. We then extend this technique to the multiregion level to determine the contribution of interactions between brain areas in reward-related decision making.

To address this, subjects performed a probabilistic reversal task (4, 15) while being scanned with fMRI. On each trial, subjects are presented with two fractal stimuli and asked to select one (Fig. 1A), with the objective of accumulating as much money as possible. After making a choice, subjects receive either a monetary gain or a monetary loss. However, one choice is "correct" in that choosing that stimulus leads to a greater probability of winning money and hence to an accumulating monetary gain, whereas the other choice is "incorrect" in that choosing that stimulus leads to a greater probability of losing money and hence to an accumulating monetary loss. After a time, the contingencies reverse so that what was the correct choice becomes the incorrect choice and vice versa. To choose optimally, subjects need to work out which stimulus is correct and continue to choose that stimulus until they determine that the contingencies have reversed, in which case they should switch their choice of stimulus. The goal of our study is to decode subjects' behavioral choices on a subsequent trial on the basis of neural activity on the preceding trial.

An important feature of this task is that its probabilistic nature precludes subjects from inferring which stimulus is correct on the basis of the outcome received on the previous trial alone, because both correct and incorrect stimulus choices are associated with rewarding and punishing feedback. Rather, subjects need to take into account the history of outcomes received to make decisions about what choices to make in future. Further-

NEUROSCIENCE

A Trial design

BOLD response

Onset 2s Fix. 4s Fixation 12s Trial ends

~.5s Subject
makes choice
3s Reward
delivered

Three scans used for
decision decoding

B Region classifier construction

within a region

Between regions

MPFC
mOFC
lOFC
Amygdala
vStriatum
dStriatum
DLPFC
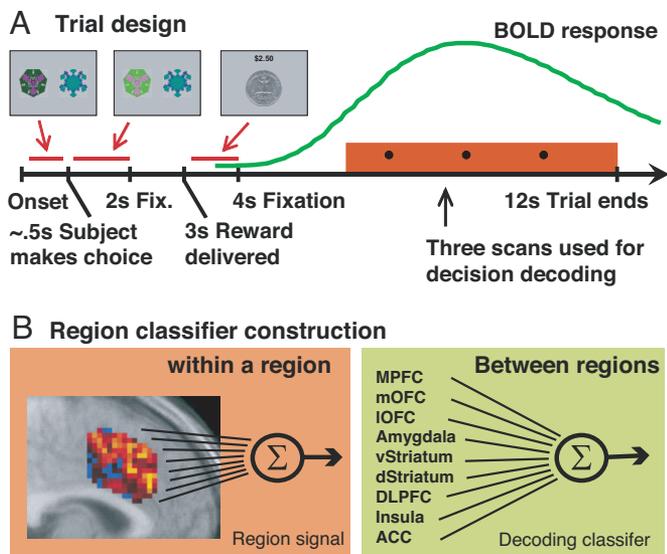Insula
ACC

Region signal   Decoding classifier

Fig. 1. Task outline and classifier construction. (*A*) Reversal task setup. Subjects chose one of two fractals, which on each trial were randomly placed to the left or right of the fixation cross. The chosen stimulus is illuminated until 2 s after the trial onset. After a further 1 s, a reward (winning 25 cents) or punishment (losing 25 cents) is delivered for 1 s, with the total money earned displayed at the top. The screen is then cleared, and a central fixation cross is presented for 8 s before the next trial begins. One stimulus is designated the correct stimulus, in that choosing that stimulus leads to a monetary reward on 70% of occasions and a monetary loss 30% of the time. The other stimulus is ''incorrect,'' in that choosing that stimulus leads to a reward 40% of the time and a punishment 60% of the time. After subjects choose the correct stimulus on four consecutive occasions, the contingencies reverse with a probability of 0.25 on each successive trial. Subjects have to infer that the reversal took place and switch their choice, at which point the process is repeated. The last three scans in a trial are used by our classifier to decode whether subjects will switch their choice or not in the next trial. A canonical BOLD response elicited at the time of reward receipt is shown (in green) to illustrate the time points in the trial at which the hemodynamic response is sampled for decoding purposes. A new trial was triggered every 12 s to ensure adequate separation of hemodynamic signals related to choices on consecutive trials. The average of three scans between the outcome of reward and the time of choice in the next trial was used for decoding subjects' behavioral choice in the next trial. These three time points will not only contain activity from the decision itself (activity taking place after the receipt of feedback, but before the next trial) but also activity from the reward/punishment received in the current trial and activity consequent to the choice made in the current trial. (*B*) The multivariate region classifier used in this study is divided in two parts. The first extracts a representative signal from each region of interest (*Left*) by averaging the brain voxels within a region weighted by the voxels' discriminability of the switch vs. stay conditions. To avoid overfitting the fMRI data, we did not take into consideration the correlations between voxels within a region of interest (Eq. 3). The second part of the classifier (*Right*) adds up the signal from each region, weighted by the region's importance in classifying the subject's decision (Eq. 2). Weights are calculated by using a multivariate classifier that uses each region's decoding strength, and correlations between regions, to maximize the accuracy of the classifier in decoding whether subjects are going to switch or stay (see *Discriminative Analysis*).

more, the two stimuli are presented at random on the left or right of the screen. Thus, on the previous trial, subjects do not know in advance which of two possible motor responses are needed to implement a particular decision until such time as the next trial is triggered. Consequently, our fMRI signal cannot be driven merely by trivial (i.e., non-decision-related) neural activity pertaining to preparation of a specific motor response (choose left vs. choose right) because such signals are not present before the stimuli are shown. Therefore, the only signals in the brain relevant to decoding choice are those pertaining to the subjects' abstract decision of whether to maintain their current choice of

stimulus or switch their choice to the alternative stimulus, or else to those pertaining to the consequences of that decision (e.g., to implement a switch in response set).

Nine regions of interest were specified *a priori* (see *Materials and Methods*), based on previous literature implicating these regions in reward-related decision making. These include the medial and lateral orbitofrontal cortex and adjacent mPFC. These regions have been shown to encode expected reward values, as well as the reward value of outcomes (4, 16, 17). Moreover, signals in these regions have been found to relate to behavioral choice, whereby activity increases in mPFC on trials when subjects maintain their current choices on subsequent trials compared with when they switch (14).

Another region that we hypothesized might contain signals relevant to behavioral choice is the ACC. This area is engaged when subjects switch their choice of stimulus on reversal learning tasks (13, 14), suggesting that signals there relate to behavioral choice. A general role for this region in monitoring action-outcome associations has recently been proposed (18). The region has also been argued to mediate action selection under situations involving conflict between competing responses (7) and action selection between responses with different reward contingencies (19). ACC has also been suggested to play a role in monitoring errors in behavioral responding or even in decoding when these errors might occur (20). What all of these accounts of anterior cingulate function have in common is that they posit an intervening role for this area between the processing of a stimulus input and the generation of an appropriate behavioral response, even though such accounts differ as to precisely how this region contributes at this intermediate stage. On these grounds, we hypothesized that neural signals in ACC would be relevant for decoding subsequent behavioral choices.

Other regions we deemed relevant to decision making include the insular cortex, which has been shown to respond during uncertainty in action choice, as well as under situations involving risk or ambiguity (21–23). Kuhnen and Knutson (24) showed that neural activity in this region on a previous trial correlated with whether subjects will make a risk-seeking or risk-averse choice in a risky decision-making paradigm. We also include in our regions of interest ventral striatum, where activity is linked to errors in prediction of future reward, and dorsal striatum, which is argued to mediate stimulus-response learning and goal-directed action selection (25–28). Another region we included is the amygdala, which has been implicated in learning of stimulus-reward or stimulus-punisher associations (29–31).

We analyzed the contribution each region of interest gives to the decoding of choice behavior in two ways. In the first, we study each region individually and compare their discriminative power for decoding behavioral choice. This is done by separating fMRI signals in each region into spatially local and spatially global signals, thus disambiguating results that correspond to classic fMRI approaches (global signals), with results that can only be obtained by using multivariate fMRI decoding techniques (local signals). In the second approach, we make use of neural responses in all of our nine regions of interest to decode behavioral decisions by using a multivariate analysis that optimally combines information from the different brain regions (Fig. 1*B*). This method enables us to obtain better decoding accuracy than by using each region separately and to explore the relative contributions of each of these different areas to the final behavioral choice.

Results

Local vs. Global Signals Related to Behavioral Choice in Regions of Interest. To address the contribution of global vs. local signals in the encoding of behavioral choice, we separated the original fMRI data into global signals with a spatial scale bigger than 8 mm and local signals with a spatial scale smaller than 8 mm (see

Fig. 2. Global and local fMRI signals related to behavioral choice. (*A*) Here, we show fMRI signals related to behavioral choice, i.e., whether subjects will switch or maintain (stay) their choices on a subsequent trial. Voxel *t*-scores for the discriminability between switch and stay trials is shown for two individual subjects with data in its original form (*Left*) and then decomposed into a global spatial component (with spatial scale >8 mm; *Center*) and a local spatial component (spatial scale <8 mm; *Right*). The ACC region of interest is outlined in white for reference. Red and yellow indicate increased responses on switch compared with stay trials, whereas blue colors indicate stronger responses on stay compared with switch trials (also see SI Fig. 10). (*B*) Results from a group random effects analysis across subjects conducted separately for the original unsmoothed data, global data, and local data. Whereas global signals survive at the random effects level (consistent with classical fMRI analyses), local spatial signals do not survive at the group random effects level. Random effect *t*-scores are shown with a threshold set at $P < 0.2$ for visualization.

Fig. 3. Illustration of the decoding accuracy for subjects' subsequent behavioral choices for each individual region and combination across regions. (*A*) Plot of average accuracies across subjects shown separately for local and global spatial scales. Both spatial scales contain information that can be used to decode subjects' subsequent behavioral choice, in all of our regions of interest. Notably, decoding accuracies are comparable at the local and global scales within each region. (*B*) Plot of average across accuracy across subjects for each region individually combining both local and global signals. (*C*) Results of the hierarchical multiregion classifier analysis, averaged across subjects. An ordering of regions was performed by starting with a classifier that only contains the individual region with best overall accuracy (ACC; leftmost column), and iteratively adding to this classifier the regions whose inclusion increases the accuracy of the classifier the most (or decreases the least). Thus, the second column shows the accuracy of a classifier containing ACC and ventral striatum, the third column the accuracy of a classifier containing ACC, ventral striatum, and mPFC, and so forth. The combination of the three regions that provide the best decoding accuracy are highlighted in gray. Addition of a fourth region (dorsal striatum) does not significantly increase decoding accuracy. All error bars indicate standard errors of the mean. (*D*) Decoding accuracy for the three region classifier shown separately for each individual subject (also see SI Table 1).

Materials and Methods). Fig. 2*A* shows the statistical significance of each voxel when discriminating between switch vs. stay decisions in two subjects. Local signals (Fig. 2*A Right*) defined this way do not survive classical fMRI analysis (Fig. 2*B*) and can only be studied by using signal analysis techniques sensitive to spatially distributed signals. We evaluated the degree to which each individual region of interest could decode subjects' subsequent choices when using either global or local signals (Fig. 3*A*). Each subject underwent four separate fMRI sessions (70 trials each) during which they performed the decision-making task. Four classifiers were trained and tested for each subject by using four-fold cross validation, where each classifier is trained by using three of the sessions (210 trials) and then tested on the session that is left out (70 trials). Decoding accuracy derived from global and local signals was comparable within each region of interest, suggesting that local and global signals strongly covary in each of the regions studied. There was a trend toward a greater contribution of local signals compared with global signals in overall decoding accuracy in ACC, although this did not reach statistical significance (at $P < 0.08$).

Decoding Accuracy of Each Individual Region. When combining both local and global signals and evaluating decoding accuracy for each region alone, we find that each region can decode better than chance whether a subject is going to switch or not (Fig. 3*B*), with the highest accuracies being obtained by ACC (64%), anterior insula (62%), and DLPFC (60%). To address whether

the difference in decoding accuracy across regions is merely a product of intrinsic differences in MR signal to noise in these areas, we examined the signal-to-noise ratio in each region by analyzing responses elicited by the main effect of receiving an outcome compared with rest. All regions had comparable signal-to-noise ratio to the main effect of outcome receipt [supporting information (SI) Fig. 4], suggesting that accuracy differences are unlikely to be accounted for by variations in intrinsic noise levels between regions.

Combined Accuracy Across Multiple Regions. Next, we aimed to determine whether a combination of specific brain regions would provide better decoding accuracy than when just considering one region alone. For this, we built regional classifiers using a multivariate approach that takes into account interactions between multiple regions of interest when decoding decisions (see Fig. 1*B*). This approach optimally combines both local and global signals from each region of interest. To determine which subset of regions to include in our classifier, we performed a hierarchical analysis whereby we started with the most accurate individual region and then iteratively built multiregion classifiers by adding one region at a time. At each step in this iterative process, we added the region that increased the multiregion

NEUROSCIENCE

classifier's accuracy the most (out of the remaining regions). Fig. 3C shows the results of this process. We found that out of our nine regions of interest, a classifier with only three of these areas (ACC, mPFC, and ventral striatum) achieved an overall decoding accuracy of $67 \pm 2\%$, a significantly better decoding accuracy of subject's choice than that provided by each region alone (for example, compared with ACC at $P < 0.01$). Accuracy increase when adding regions is not only due to the signals related to behavioral choice in each region but also depends on the degree of statistical independence of noise across regions (SI Fig. 5). Fig. 3D shows the average accuracy for each individual subject when using our region-based classifier. Receiver operating characteristic curves representing the average classifier accuracy across a range of response thresholds are shown in SI Fig. 6 (see also SI Table 2).

Insula and DLPFC, which on their own have high decoding accuracy, were not selected in our hierarchical classifier, suggesting that signals from these regions are better accounted for by the other included regions. To account for the possibility that another combination of regions could substitute equally well for the regions included in the hierarchical classifier, we ran an additional analysis whereby we tested the classification accuracy of every possible combination of three regions (SI Fig. 7). Even in this case, we still found that the specific combination of regions identified from the hierarchical analysis were highest in decoding accuracy compared with all other possible combinations, supporting the conclusion that the specific regions we identified are sufficient for decoding decisions up to the overall decoding accuracy obtained in our study.

It should be noted that the approach we use here whereby signals are combined across regions proved significantly better at decoding behavioral decision making than alternative decoding techniques that do not employ this multiregion approach (SI Fig. 8). However, when in the hierarchical analysis we added more than four regions, the combined classifier's accuracy gradually decreased again (Fig. 3C), perhaps because of over-fitting of the training data.

Decisions per Se or Detection of Rewarding vs. Punishing Outcomes?
A key question is whether the decoding accuracy of our regional classifier is derived by detecting activity elicited by the decision process and its consequences or merely reflects detection of the sensory and affective consequences of receiving a rewarding or punishing outcome on the preceding trial. To test this, we restricted input to the classifier to only those trials on which subjects received a punishing outcome. Even in this instance, the classifier was able to decode subjects' decisions to switch or stay on the subsequent trial with $57 \pm 1\%$ accuracy, significantly better than chance (at $P < 10^{-8}$, across-subjects mean accuracy). This finding suggests that our classifier is using information relating to the behavioral decision itself and is not merely discriminating between rewarding or punishing outcomes on the immediately preceding trial. Additional analyses in support of this conclusion are detailed in *SI Methods*.

Discussion
The results of this study demonstrate that it is possible to decode, with a high degree of accuracy, reward-related decisions or the consequence of those decisions (in terms of initiating a change in response choice) in human subjects on the basis of neural activity measured with fMRI before the specific physical action involved is either planned or executed. Theoretical accounts of goal-directed behavior differ in the degree to which decisions are suggested to be linked to the specific action needed to carry them out. According to one stimulus-driven view, decisions are computed abstractly in terms of the specific stimulus or goal the subject would like to attain (32, 33). In the case of the probabilistic reversal task used here, this comes down to a choice

between which of two different fractal stimuli to select. An alternative approach is to propose that decisions are computed by choosing between the set of available physical actions that are required to attain a particular goal. Here, we measure neural responses on a preceding trial before subjects are presented with the explicit choice between two possible actions, and before subjects know which specific action they will need to select to implement that decision. The fact that these decisions can be decoded before subjects are aware of the specific action that needs to be performed to realize them (choose left or right button) suggests that decision signals can be encoded in the brain at an abstract level, independently of the actual physical action with which they are ultimately linked (32). We should note, however, that our decoding technique, which is based on activity elicited at the time of receipt of the outcome on the preceding trial, is likely to be picking up both the decision itself and the consequence of the decision. In other words, once a decision to switch is computed, a change in stimulus-response mapping is going to be initiated, and the activity being detected in our analysis may also reflect this additional process.

Our findings have important implications not only for understanding what types of decisions are computed but also when these decisions are computed. In the context of the reversal task, it is possible for a decision to be computed at any point in time between receipt of the outcome on the previous trial and implementation of the behavioral choice on the next. By using multivariate fMRI techniques, we have been able to show that subsequent decisions (or the consequences of those decisions) can be decoded on the basis of signals present on the preceding trials (after outcomes are received). This suggests that the decision to switch or maintain current response set may be initiated as soon as the information needed to compute the decision is available, rather than being implemented only when required on the subsequent trial.

In this study, we also separated fMRI signals with a global spatial encoding from those with a local spatial encoding and evaluated the information each contained for decoding behavioral choice within each of our regions of interest. Global signals, as we define them here, are relatively uniform spatially extended clusters of activation within a given area (with a spatial scale of >8 mm). Typically, these spatially smoothed signals are those reported in conventional fMRI analyses, because they are especially likely to survive at the group level. However, recently it has been shown that information about task processes can be obtained from considering local spatially distributed variations in voxel activity (9, 10). In the present case, we defined spatially local signals as those with a spatial scale of <8 mm. In this study, we showed that within each region of interest, local signals do convey important information regarding behavioral choice over and above that conveyed by the global signals. However, we did not find strong evidence for a dissociation between regions in the degree to which they were involved in encoding local and global signals, except for a trend in ACC toward a greater role for this region in encoding local as opposed to global signals. These results suggests that at least in the context of the present reversal learning task, the presence of global and local information relevant to behavioral decision making strongly covaries within areas. This is in contrast to results observed in the visual system, where in some instances local signals convey information pertaining to visual perception even when global signals do not. Local fMRI signals in visual cortex have been argued to relate to the columnar organization in this area of the brain. It should be noted, however, that much less is known about the degree to which columnar organization exists outside of visual cortical areas, and hence, the underlying neural architecture that contributes to local fMRI signals in other areas of the brain such as the prefrontal cortex remains to be understood.

We also used a multivariate analysis technique whereby the degree to which neural signals in multiple brain regions contribute to the decision process are evaluated simultaneously. This approach has allowed us to efficiently recruit signals from diverse brain regions to arrive at a better decoding accuracy for subjects' behavioral choice than would follow from considering activity in any one region alone. As a consequence, our findings suggest that reward-related decision processes might be better understood as a product of computations performed across a distributed network of brain regions, rather than being the purview of any one single brain area.

Nevertheless, our results do suggest that some regions are more important than others. When we compared the decoding accuracy of classifiers incorporating information from all of our regions of interest to the accuracy of classifiers using information derived from different subsets of regions, we found that activity in a specific subset of our regions of interest accounted for the maximum accuracy of our classifier; namely, the ACC, mPFC, and ventral striatum. Each of these regions has been identified previously as playing a role in decision making and behavioral choice on the basis of prior fMRI studies using traditional statistical analysis techniques (13, 14, 27). Out of these, one region in particular stood out as contributing the most: dorsal ACC. This region has previously implicated in diverse cognitive functions, including response conflict and error detection (7, 20). However, a recent theoretical account has proposed a more general role for this region in guiding action selection for reward (13, 18, 19). Although not incompatible with response-conflict or error-detection theories, our results are especially consistent with this latter hypothesis, suggesting that this region is playing a key role in implementing the behavioral decision itself.

Some of the regions featured in this study, such as DLPFC, may contribute to task or cognitive-set switching more generally (34–36) and are unlikely to be uniquely involved in reward-related decision making. However, it is notable that DLPFC was ultimately not selected in our combined classifier. Instead, the regions that were selected have previously been specifically implicated in reward-related learning and/or in implementing changes in behavior as a consequence of such learning and not in cognitive set-shifting *per se* (13, 14).

The present study demonstrates that it is possible to decode, on a single-trial basis, abstract reward-related decisions in human subjects, in essence, by reading their decision before the action is executed. Our findings are consistent with the proposal that decision making is best thought of as an emergent property of interactions between a distributed network of brain areas rather than being computed in any one single brain region. Of all of the regions we studied, we found that a subset of three regions seemed to contain information that was sufficient to decode behavioral decision making: ACC, mPFC, and ventral striatum. Future studies are needed to determine whether these regions contain information specifically required for probabilistic reversal learning or whether other types of reward-related decisions can also be decoded on the basis of information contained in these areas.

Materials and Methods

Subjects. Eight healthy right-handed normal subjects participated in this study (four female, mean age 27.6 ± 5.6 years). The subjects were preassessed to exclude those with a prior history of neurological or psychiatric illness. All subjects gave informed consent, and the study was approved by the Institute Review Board at the California Institute of Technology. Subjects were paid according to their performance in the task. Before scanning, subjects were trained on three different versions of the probabilistic reversal task, as in Hampton *et al.* (17), as described in *SI Methods*.

Data Acquisition and Preprocessing. The functional imaging was conducted by using a Siemens 3.0-T Trio MRI scanner to acquire gradient echo T2* weighted echo-planar images with blood oxygenation level-dependent (BOLD) contrast. To optimize functional sensitivity in orbitofrontal cortex, we used a tilted acquisition in an oblique orientation of 30° to the anterior–posterior commissure line. Four sessions of 450 volumes each (4 × 15 min) were collected in an interleaved-ascending manner. The imaging parameters were as follows: echo time, 30 ms; field-of-view, 192 mm; in-plane resolution and slice thickness, 3 mm; TR, 2 s. Whole-brain high-resolution T1-weighted structural scans (1 × 1 × 1 mm) were acquired from each subject and co-registered with their mean echo-planar image. Image analysis was performed by using SPM2 (Wellcome Department of Imaging Neuroscience, Institute of Neurology, London, U.K.). To correct for subject motion, the images were realigned to the first volume, spatially normalized to a standard T2* template with a resampled voxel size of 3 mm. Trials were 12 s long and were time locked to the start of the fMRI echo-planar scan sequence. This was done to ensure that the scans from the previous trial used to decode the subject's decision in the next trial would not be contaminated with BOLD activity arising from the choice itself on the subsequent trial. A running high-pass filter (the mean BOLD activity in the last 36 volumes, or 72 s, was subtracted from the activity of the current volume) was also applied to the data. This was used instead of the usual high-pass filtering (37) so that BOLD activity in a volume would not be contaminated with activity from the choice itself in subsequent volumes. In the scanner, visual input was provided with Restech (Resonance Technologies, Northridge, CA) goggles, and subjects used a button box to choose a stimulus.

Global and Local Spatial Signals. To dissociate global and local signals relevant to behavioral choice, we used the following procedure. (*i*) The activity in each voxel was scaled such that the variance of the BOLD activity over all trials in a session was equalized across all voxels. (*ii*) The fMRI data were spatially smoothed by using a Gaussian kernel with a full width at half maximum of 8 mm, to capture global changes in signal. (*iii*) fMRI data containing only locally distributed spatial signals were then extracted by subtracting the smoothed fMRI data (obtained in *ii*), from the non-spatially smoothed fMRI data (obtained in *i*). This procedure adopts the assumption that BOLD activity is a function of the underlying neuronal activity that is identical across neighboring voxels, except for a scaling constant. Furthermore, *i* estimates and eliminates the scaling differences across voxels, but errors in the estimation of this scaling could lead to an incomplete dissociation between local and global signals. The procedure also assumes that if local encodings exist, they will have the same scaling characteristics across all brain regions.

Region-of-Interest Specification. Nine regions of interest were specified based on previous literature implicating these regions in reward-related decision making, and delineated by anatomical landmarks (SI Fig. 9). Regions of interest were specified by using a series of spheres centered at specified (*x*, *y*, *z*) Montreal Neurological Institute coordinates and with specified radii in millimeters (see SI Table 3 for complete specification).

Discriminative Analysis. To optimally classify whether subjects will switch or stay in a given trial, the fMRI voxel activity **x** (see Fig. 1*A*) is assigned to the action a_i for which the posterior probability $p(a_i|\mathbf{x}) = p(\mathbf{x}|a_i)p(a_i)/p(\mathbf{x})$ is maximal. Here, $p(\mathbf{x}|a_i)$ is the distribution of voxel activities given action a_i. Assuming that the fMRI activity **x** follows a multivariate normal distribution with the same covariance matrix Σ given either action, the posterior probability whether to choose the switch action is

$$p(a_{\mathrm{sw}}|\mathbf{x})$$

$$= \frac{p(\mathrm{sw})e^{-\frac{1}{2}(\mathbf{x} - \bar{\mathbf{x}}_{\mathrm{sw}})^T S^{-1}(\mathbf{x} - \bar{\mathbf{x}}_{\mathrm{sw}})}}{p(\mathrm{sw})e^{-\frac{1}{2}(\mathbf{x} - \bar{\mathbf{x}}_{\mathrm{sw}})^T S^{-1}(\mathbf{x} - \bar{\mathbf{x}}_{\mathrm{sw}})} + p(\mathrm{st})e^{-\frac{1}{2}(\mathbf{x} - \bar{\mathbf{x}}_{\mathrm{st}})^T S^{-1}(\mathbf{x} - \bar{\mathbf{x}}_{\mathrm{st}})}},$$

$$[1]$$

where $\bar{\mathbf{x}}_{\mathrm{sw}}$ and $\bar{\mathbf{x}}_{\mathrm{st}}$ are the training sample means of the fMRI activity for the switch and stay actions, respectively, and S is the pooled covariance matrix estimated from the training sample. This can be simplified to $p(a_{\mathrm{sw}}|\mathbf{x}) = 1/(1 + e^{-y})$, where

$$y = \mathbf{w}^T\mathbf{x} - \theta,$$
$$\mathbf{w}^T = (\bar{\mathbf{x}}_{\mathrm{sw}} - \bar{\mathbf{x}}_{\mathrm{st}})^T S^{-1}. \qquad [2]$$

Here, θ is a threshold variable that groups all constants. Given the brain voxel activity $\mathbf{x}$ in a single trial, choosing the action with maximal posterior is equivalent to choosing the action for which $y > 0$.

The classifier was built in two steps (see Fig. 1B). In the first, nine regions of interest were specified (SI Fig. 9), and a unique signal from each region was obtained by adding up the activities of all voxels in that region, weighted by the voxels' discriminability:

$$r = \mathbf{w}^T\mathbf{x},$$
$$w_i = \frac{\bar{x}^i_{\mathrm{sw}} - \bar{x}^i_{\mathrm{st}}}{\sigma_i^2}, \qquad [3]$$

where σ_i^2 is the pooled voxel variance of voxel i. This approach assumes that the noise is independent across voxels, a procedure used to avoid overfitting the classifier to the fMRI data. The second step utilizes the coalesced regional activities as input to a full Gaussian discriminative classifier (Eq. **2**), where weights are assigned to each region.

Decoding accuracy is measured as the percentage of correctly decoded behavioral choices. That is, the mean between correctly decoded switch actions (number of correctly decoded switches divided by the total number of switch actions) and correctly decoded stay actions (number of correctly decoded stays divided by the total number of stay actions). This measure takes into account the fact that the number of times a subject switches or stays can be different across sessions.

1. Platt ML, Glimcher PW (1999) *Nature* 400:233–238.
2. Sugrue LP, Corrado GS, Newsome WT (2004) *Science* 304:1782–1787.
3. Schultz W, Dayan P, Montague PR (1997) *Science* 275:1593–1599.
4. O'Doherty J, Kringelbach ML, Rolls ET, Hornak J, Andrews C (2001) *Nat Neurosci* 4:95–102.
5. Cohen JD, Botvinick M, Carter CS (2000) *Nat Neurosci* 3:421–423.
6. Gehring WJ, Knight RT (2000) *Nat Neurosci* 3:516–520.
7. Kerns JG, Cohen JD, MacDonald AW, III, Cho RY, Stenger VA, Carter CS (2004) *Science* 303:1023–1026.
8. Haxby JV, Gobbini MI, Furey ML, Ishai A, Schouten JL, Pietrini P (2001) *Science* 293:2425–2430.
9. Kamitani Y, Tong F (2005) *Nat Neurosci* 8:679–685.
10. Haynes JD, Rees G (2005) *Nat Neurosci* 8:686–691.
11. Polyn SM, Natu VS, Cohen JD, Norman KA (2005) *Science* 310:1963–1966.
12. Pessoa L, Padmala S (2005) *Proc Natl Acad Sci USA* 102:5612–5617.
13. Bush G, Vogt BA, Holmes J, Dale AM, Greve D, Jenike MA, Rosen BR (2002) *Proc Natl Acad Sci USA* 99:523–528.
14. O'Doherty J, Critchley H, Deichmann R, Dolan RJ (2003) *J Neurosci* 23:7931–7939.
15. Cools R, Clark L, Owen AM, Robbins TW (2002) *J Neurosci* 22:4563–4567.
16. Knutson B, Taylor J, Kaufman M, Peterson R, Glover G (2005) *J Neurosci* 25:4806–4812.
17. Hampton AN, Bossaerts P, O'Doherty JP (2006) *J Neurosci* 26:8360–8367.
18. Walton ME, Devlin JT, Rushworth MFS (2004) *Nat Neurosci* 7:1259–1265.
19. Williams ZM, Bush G, Rauch SL, Cosgrove GR, Eskandar EN (2004) *Nat Neurosci* 7:1370–1375.
20. Brown JW, Braver TS (2005) *Science* 307:1118–1121.
21. Critchley HD, Mathias CJ, Dolan RJ (2001) *Neuron* 29:537–545.
22. Hsu M, Bhatt M, Adolphs R, Tranel D, Camerer CF (2005) *Science* 310:1680–1683.
23. Huettel SA, Stowe CJ, Gordon EM, Warner BT, Platt ML (2006) *Neuron* 49:765–775.
24. Kuhnen CM, Knutson B (2005) *Neuron* 47:763–770.
25. Delgado MR, Locke HM, Stenger VA, Fiez JA (2003) *Cognit Affect Behav Neurosci* 3:27.
26. O'Doherty JP, Dayan P, Friston K, Critchley H, Dolan RJ (2003) *Neuron* 38:329–337.
27. O'Doherty J, Dayan P, Schultz J, Deichmann R, Friston K, Dolan RJ (2004) *Science* 304:452–454.
28. Yin HH, Ostlund SB, Knowlton BJ, Balleine BW (2005) *Eur J Neurosci* 22:513–523.
29. Gottfried JA, O'Doherty J, Dolan RJ (2003) *Science* 301:1104–1107.
30. Schoenbaum G, Setlow B, Saddoris MP, Gallagher M (2003) *Neuron* 39:855–867.
31. Holland PC, Gallagher M (2004) *Curr Opin Neurobiol* 14:148–155.
32. Padoa-Schioppa C, Assad JA (2006) *Nature* 441:223–226.
33. Rolls ET (2000) *Cereb Cortex* 10:284–294.
34. MacDonald AW, Cohen JD, Stenger VA, Carter CS (2000) *Science* 288:1835–1838.
35. Brass M, von Cramon DY (2004) *J Neurosci* 24:8847–8852.
36. Dias R, Robbins TW, Roberts AC (1996) *Nature* 380:69–72.
37. Friston KJ, Ashburner J, Frith CD, Poline JB, Heather JD, Frackowiak RSJ (1995) *Hum Brain Mapp* 3:165–189.