NeuroImage 70 (2013) 66–79

Contents lists available at SciVerse ScienceDirect

NeuroImage

journal homepage: www.elsevier.com/locate/ynimg

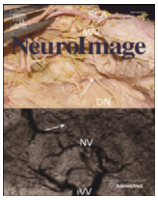

Action selection in multi-effector decision making

Seth Madlon-Kay [a,*], Bijan Pesaran [b], Nathaniel D. Daw [b,c]

[a] Center for Cognitive Neuroscience, Duke University, Durham, NC 27708, USA
[b] Center for Neural Science, New York University, New York, NY 10003, USA
[c] Department of Psychology, New York University, New York, NY 10003, USA

ARTICLE INFO

Article history:
Accepted 2 December 2012
Available online 7 December 2012

Keywords:
Bimanual
Decision making
Value maps
Effector specificity
vmPFC

ABSTRACT

Decision making and reinforcement learning over movements suffer from the curse of dimensionality: the space of possible movements is too vast to search or even represent in its entirety. When actions involve only a single effector, this problem can be ameliorated by considering that effector separately; accordingly, the brain's sensorimotor systems can subdivide choice by representing values and actions separately for each effector. However, for many actions, such as playing the piano, the value of an action by an effector (e.g., a hand) depends inseparably on the actions of other effectors. By definition, the values of such coordinated multi-effector actions cannot be represented by effector-specific action values, such as those that have been most extensively investigated in parietal and premotor regions. For such actions, one possible solution is to choose according to more abstract valuations over different goods or options, which can then be mapped onto the necessary motor actions. Such an approach separates the learning and decision problem, which will often be lower-dimensional than the space of possible movements, from the multi-effector movement planning problem. The ventromedial prefrontal cortex (vmPFC) is thought to contain goods-based value signals, so we hypothesized that this region might preferentially drive multi-effector action selection.

To examine how the brain solves this problem, we used fMRI to compare patterns of BOLD activity in humans during reward learning tasks in which options were selected through either unimanual or bimanual actions, and in which the response requirements in the latter condition inseparably coupled valuation across both hands. We found value signals in the bilateral medial motor cortex and vmPFC, and consistent with previous studies, the medial motor value signals contained contra-lateral biases indicating effector-specificity, while the vmPFC value signals did not exhibit detectable effector specificity. Although neither region's value signaling differed significantly between bimanual and unimanual conditions, the vmPFC value region showed greater connectivity with the medial motor cortex during bimanual than during unimanual choices. The specific region implicated, the anterior mid-cingulate cortex, is thought to act as a hub that links subjective value signals to motor control centers. These results are consistent with the idea that while valuation for unilateral actions may be subserved by an effector-specific network, complex multi-effector actions preferentially implicate communication between motor regions and prefrontal regions, which may reflect increased top-down input into motor regions to guide action selection.

Introduction

Humans and other animals possess a variety of effectors, which together support a wide array of possible motor actions. The large space of possible movements implies a classic “curse of dimensionality” due to the exponential explosion of combinations of movements. For instance, coarsely considering only actions involving the eyes, arms, hands and legs yields a seven-dimensional space that can contain many millions of possible joint actions, even if each effector considered alone has only a few candidate movements. Given that the brain most certainly parses the body more finely than merely seven effectors, the actual number of possibilities must be many orders of magnitude larger. Representing such a large number of actions, let alone deciding between them, is computationally challenging. The problem becomes especially laborious when appropriate decisions are learned by trial and error, since even trying out all the possibilities is clearly impossible.

One method by which the brain seems to cope with this curse is by subdividing decisions into lower dimensional, effector-specific subspaces: that is, controlling particular effectors separately and independently from the others. Indeed, the basic organization of the brain seems well suited to a “divide and conquer” strategy. Primary motor cortex and many affiliated territories including basal ganglia, supplementary and premotor areas, contain effector-specific topographies wherein different muscle groups in different parts of the body are governed by spatially segregated populations of neurons (Alexander et al., 1986; Blanke et al., 2000; Leyton and Sherrington, 1917; Pesaran et al., 2006; Schlag

---

* Corresponding author at: 6 Washington Place rm 809, New York, NY 10003, USA.
E-mail addresses: seth.madlonkay@duke.edu (S. Madlon-Kay), bijan@nyu.edu (B. Pesaran), nathaniel.daw@nyu.edu (N.D. Daw).

1053-8119/$ – see front matter © 2012 Elsevier Inc. All rights reserved.
http://dx.doi.org/10.1016/j.neuroimage.2012.12.001

and Schlag-Rey, 1987). Similarly effector-specific systems can be found in sensorimotor areas such as posterior parietal cortex, where saccadic eye movements and reaches are represented by distinct subregions (Batista et al., 1999; Calton et al., 2002; Levy et al., 2007; Snyder et al., 1997). Recent evidence suggests that such effector-specific movement representations are also a substrate for evaluating and choosing among these actions, e.g. with the eye movement representations in parietal cortex carrying 'value maps' of rewards expected for different saccades and with action values for movements of the hands and eyes appearing in distinct subregions of supplementary motor cortex (Andersen and Buneo, 2002; Palminteri et al., 2009; Platt and Glimcher, 1999; Scherberger and Andersen, 2007; Shadlen and Newsome, 2001; Sugrue and Corrado, 2004; Wunderlich et al., 2009). Furthermore, in a recent fMRI study we tested explicitly for subdivision across effectors (here, left and right hands) and found lateralized patterns of BOLD signals related to valuation and reward prediction errors consistent with the separation of a multi-effector task into effector-specific subproblems (Gershman et al., 2009).

However, separating the control of effectors in this way introduces a related problem of coordination. Many actions require the conjoint action of multiple effectors working together to accomplish a single goal. When playing piano chords, for example, playing a C with the left hand will sound dissonant if the right hand chooses a B but pleasingly consonant if instead the right hand plays E. In these cases, the value of a particular effector's movements depends on the action of another effector. This means that the fundamental units of evaluation and decision must be combinations of movements spanning multiple effectors, and cannot be represented by separate maps of each effector's outcomes alone. The need for such coordination is also illustrated by corpus callosotomy patients. While they often show improved performance over neurologically intact controls in tasks in which the two hands must act completely independently, such as finger-tapping with different rhythms for each hand, they are unable to perform or learn novel actions requiring that the two hands act in concert, such as threading a needle or playing the piano (Brinkman, 1984; Franz et al., 2000; Kennerley et al., 2002). In other words, for the case of novel bimanual actions, such patients appear limited to actions that can be represented by independent, effector-specific systems.

How, then, does the brain evaluate high dimensional multi-effector actions? One approach is simply also to maintain, in the motor system, representations over multi-effector actions, and their values, as well. Indeed, single units in areas such as the supplementary motor area have been reported with differential tuning for bimanual movements relative to their unimanual counterparts (Donchin et al., 2002; Tanji et al., 1988). Such a solution would be most useful for selecting among sets of commonly-used or well-learned multi-effector actions, or within groups of effectors commonly used together. However, if applied to all possible multi-effector movements, this solution would likely run into of the curse of dimensionality because the space of all possible multi-effector actions is so vast.

Another approach, which may be complementary, is to rely on a "goods-based" rather than an "action-based" evaluation strategy (Padoa-Schioppa and Assad, 2006; Wunderlich et al., 2010): representing the value of a few possible outcomes or goods that might serve as the goals of actions, separate from the movements needed to obtain them. Once a choice has been made in the space of outcomes (which may be low-dimensional), the motor system is left to solve the more targeted problem of executing the choice. Intuitively, this is like first choosing between peanut butter and potato chips, and subsequently planning the movements required to grab the chosen snack food and extract it from its container. Such an approach would support a different sort of divide and conquer strategy, whereby learning, evaluation and choice over goals could be separated from the movement programming required to obtain them. The ventromedial prefrontal and orbitofrontal cortices (vmPFC and OFC) are thought to contain this sort of "goods-based" value signal. Neurons in the primate OFC represent values of particular consumption goods regardless of the actions involved in obtaining them, and their firing properties generally do not vary with the motor contingencies of tasks (Kennerley and Wallis, 2009; Padoa-Schioppa and Assad, 2006; Wallis and Miller, 2003). Similarly, in the ventromedial prefrontal cortex in humans, BOLD activity correlates with subjective value across a very wide and general variety of contexts, including those that do not involve making a motor action (Berns et al., 2001; Kable and Glimcher, 2007; Knutson et al., 2000; Levy and Glimcher, 2011; O'Doherty et al., 2002; Peters and Büchel, 2009; Plassmann et al., 2007; Tom et al., 2007; Wunderlich et al., 2010). The distinction between goods- and action-based choice models has been hotly debated in decision neuroscience (Glimcher, 2008; Padoa-Schioppa and Assad, 2006; Platt and Padoa-Schioppa, 2008), but these mechanisms need not be mutually exclusive, and recent imaging studies suggest that the brain uses both (Wunderlich et al., 2009, 2010). In particular, for the reasons discussed above, these two mechanisms are well suited to single- and multi-effector problems, respectively. We thus hypothesized that the vmPFC might preferentially guide action selection during multi-effector choice.

To investigate the neural mechanisms of multi-effector decision making, we used fMRI to examine patterns of neural activity in humans playing a four-armed bandit reinforcement learning task, in which choices were executed, in different conditions, by unimanual or bimanual button presses. Although in both conditions subjects learned the values of four shapes, in the bimanual condition, the response requirements for choosing a shape were such that (unlike the separable bimanual task studied by Gershman et al. (2009)) the value of each action was indivisible across the effectors. For this reason, choice could operate over effector-specific value representations in the unimanual condition, but not in the bimanual one. (Note that here we are conceptualizing each hand as a single effector; however, one could instead describe each finger as an effector and the above claims would still apply.) Using estimates of the participants' values of the options on each trial, we probed the brain for differences in patterns of value-related activity associated with bimanual versus unimanual actions. In particular, we examined functional connectivity between value-related areas of vmPFC and more posterior motor regions as a potential index of the involvement in action of abstract valuation systems.

Methods

Participants

Twenty-three right-handed adults, ages 19 to 38 (median 22, 17 females), were recruited from the New York University community to participate in the experiment in exchange for payment. Participants received a fixed amount of money for completing the experiment as well as the money they won during the task. Three participants were excluded due to technical difficulties, leaving twenty subjects whose data were used in the analysis. All participants gave informed consent, and the study was approved by the New York University Committee on Activities Involving Human Subjects.

Experimental task

Participants performed a bandit task in which they earned money based on their choices. On each trial, participants selected one of four shapes, each of which was associated with some probability of winning money (Fig. 1A). Participants received feedback regarding whether they won money on that trial 4 s after the choice was executed. The probability of winning from each shape changed slowly and independently over the course of the task according to a Gaussian random walk, with reflecting bounds below at 0.3 and above at 0.7, so maximizing payoff required continuously tracking shape values throughout the task. For each winning trial participants received $0.10 at the end of the experiment.

Participants selected shapes by pressing buttons with their index and middle fingers on their right and left hands. In order to distinguish single and multi-effector choice-related activity, we manipulated the motor actions required to select shapes. In the unimanual condition, each of the four shapes could be selected by pressing a button with one of the four fingers of the right or left hand being used in the task. In the bimanual condition, participants selected shapes by simultaneously pressing one of four possible pairs of buttons using one finger on each hand in a coordinated bimanual action (Fig. 1B). If a button was pressed with one hand and no button press was detected for the other hand within 100 ms, the trial was deemed a mistrial and no feedback was given. In both conditions, any trial in which participants did respond within 1.5 s of choice onset was also considered a mistrial, with no feedback given. To prevent participants from preparing motor action plans during the inter-trial interval (ITI) and better constrain the range of times during which motor preparation could take place, the positions of the shapes on the screen were shuffled randomly from trial to trial. Each button-pressing action thus corresponded to one of the four locations where shapes appeared on the screen rather than to a particular shape.

Participants first underwent a short training session described below, and then entered the fMRI scanner and performed the experimental task in four blocks of trials. Blocks alternated between containing only unimanual trials and containing only bimanual trials. Between blocks, participants saw text reminding them of the upcoming block type, and as an additional mnemonic, within blocks the fixation cross was red in the unimanual condition and blue in the bimanual condition. The condition of the first block was counterbalanced across participants. Each block consisted of sixty trials and lasted slightly over twelve minutes for a total of 240 trials, with 120 trials in each condition.

Training

Before entering the scanner, participants were trained on the motor actions required to select shapes in both the unimanual and bimanual conditions. The shapes appeared on the screen in the same spatial arrangement used in the main task, but with one highlighted shape appearing brighter than the rest. Participants selected the highlighted shape using the required button presses, whereupon the shapes were shuffled and a shape in a new location highlighted. Initially, cues appeared on the bottom of the screen indicating which fingers to use to select the highlighted shape, but those cues were not displayed during the second half of training. We encouraged participants to repeat the motor training if they felt they had any difficulty remembering the different action–location associations, but none performed the training more than twice. Immediately following motor training, participants played 40 practice trials of the full task, with half of the trials in each condition. All participants reported no difficulty selecting their intended shape in either condition.

Modeling choice behavior

To determine whether participants' choices were significantly guided by learning about the values of the shapes, and to generate trial-by-trial estimates of their value beliefs, we fit their choice and reward histories to a simple Rescorla–Wagner reinforcement learning (RL) model (Rescorla and Wagner, 1972), using the particular form from Schönberg et al. (2007). Under this model, on each trial t the current estimated value of the chosen shape $Q_t(s_t)$ is compared to the reward actually received that trial r_t to generate the prediction error δ_t, which is then added to $Q_t(s_t)$, weighted by learning rate α, to determine

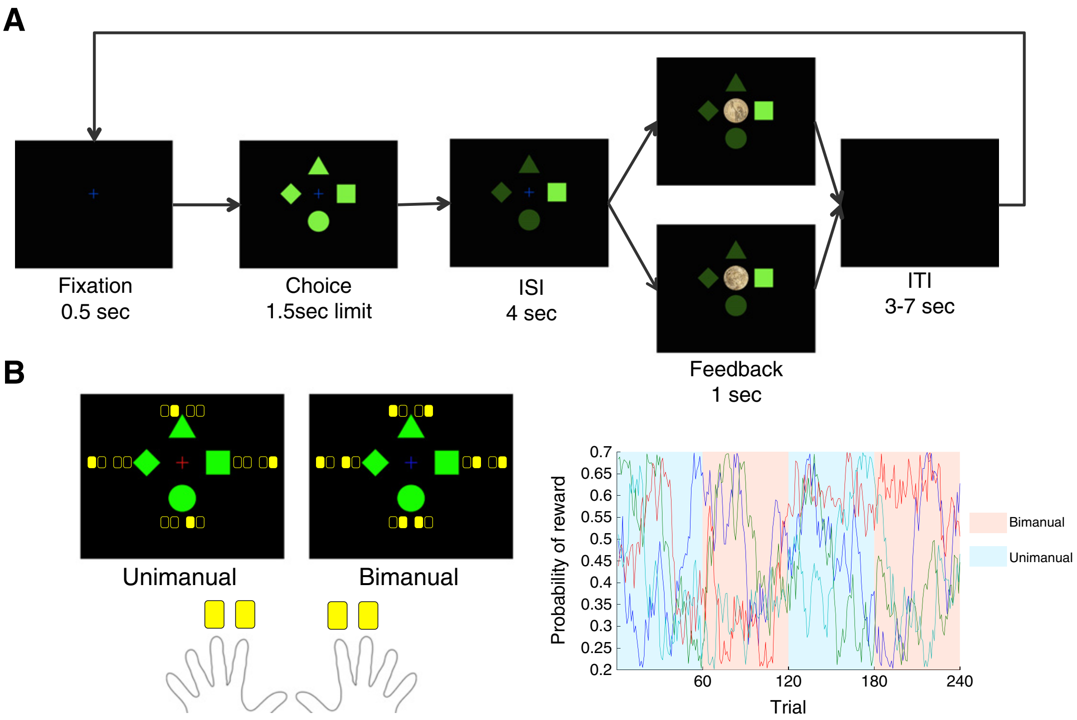

Fig. 1. Experimental task. (A) The sequence of events during a task trial. Subjects chose a shape within 1.5 s of shape onset, whereupon the unchosen shapes dimmed and, after a fixed ISI, subjects saw either a dollar coin indicating they had won money or a scrambled dollar coin indicating they had not. After a variable ITI the next trial began, with the shapes appearing in different positions. The amount of time between when a choice was made and the time limit was added to ITI to ensure consistent timing. (B) The two experimental conditions, bimanual and unimanual. Both conditions required that participants press buttons with their middle and index fingers on each hand to select shapes. A button, or combination of buttons, corresponded to a particular location on the screen and selected whatever shape occupied that location on that trial. The correspondence between buttons and choices is shown above for both conditions, with the sets of four yellow icons representing the four available buttons. The filled-in icons indicate which buttons would be used to select the neighboring shape.

$Q_{t+1}(s_t)$. This corresponds to a weighted average of past rewards from an option where the weight of a given reward decreases exponentially with samples past at a rate determined by α.

Choice history also often has strong effects on subsequent choices, with subjects often showing either perseverative tendencies or sequential sampling (Lau and Glimcher, 2005). We modeled these influences analogously to how we modeled the effects of reward, using an exponentially decaying trace κ.

Formally, the model is described by the following equations:

$$Q_{t+1}(s) = \begin{cases} Q_t(s) + \alpha * (r_t - Q_t(s)) & s = s_t \\ Q_t(s) & \forall s \neq s_t \end{cases} \tag{1}$$

$$C_{t+1}(s) = \begin{cases} C_t(s) * (1-\kappa) + \kappa & s = s_t \\ C_t(s) * (1-\kappa) & \forall s \neq s_t \end{cases} \tag{2}$$

$$\begin{aligned} V_t(s) &= \beta_Q * Q_t(s) + \beta_C * C_t(s) \\ \pi_t(s) &= \frac{exp(V_t(s))}{\sum_s \exp(V_t(s))} \end{aligned} \tag{3}$$

The free parameters α and β_Q are the learning rate and inverse softmax temperature for the effect of reward, parameters κ and β_C are the exponential decay and inverse softmax temperature associated with choice history effects, $V_t(s)$ gives the total propensity to choose shape s on trial t, and the probability of choosing s on t, $\pi_t(s)$, is derived from $V_t(s)$ via the softmax Eq. (3). Decay parameters α and κ were bounded between 0 and 1, while both inverse temperatures were unbounded. We fit the free parameters $\theta = \{\alpha, \beta_Q, \kappa, \beta_C\}$ by maximizing Eq. (4), the likelihood of the subject's series of choices D given θ with respect to θ.

$$p(D|\theta) = \prod_t \pi_t(s_t) \tag{4}$$

To ensure that learning and choice behavior were similar in both unimanual and bimanual blocks, we also fit an 'action-dependent' model where behavior was allowed to vary between blocks. This model was nearly identical to the standard RL model, differing only in that it featured two sets of parameters, θ_u and θ_b. For trials in the unimanual condition, parameter set θ_u was used to compute choice probabilities according to Eqs. (1) through (3), while trials in the bimanual condition used the θ_b parameter set.

To assess which model or mixture of models best described the population, we treated model identity as a random variable and estimated the probabilities that a random subject in the population adheres to each of the considered models (Stephan et al., 2009). We accomplished this by computing the Bayesian Information Criterion score as an approximation to the log model evidence of each model for every subject separately (Schwarz, 1978) and entered them into the function spm_BMS of the SPM8 package (Wellcome Department of Cognitive Neurology, London, UK).

Image acquisition

Functional imaging was performed on a 3T Siemens Allegra head-only scanner with a custom head coil (NM-011, Nova Medical, Wakefield, MA) located at the Center for Brain Imaging at New York University. Forty contiguous oblique-axial EPI images (3×3×3 mm voxels) were obtained each 2360 ms TR, oriented for each participant to parallel the orbital–frontal cortical surface in order to improve functional sensitivity and reduce drop-out there (Deichmann et al., 2003). Slices were positioned to ensure full coverage of orbitofrontal and prefrontal cortex, frontal motor areas, and parietal cortex. This protocol achieved nearly whole-brain coverage, though for several participants ventral parts of temporal lobe, most frequently the temporal pole, and parts of the medial temporal lobe were uncovered. A high-resolution T1-weighted anatomical image (MPRAGE sequence, 1×1×1 mm) was also acquired for each subject after the functional scans. In addition to functional and anatomical scans, we also collected for each participant a fieldmap scan (multiecho GRE sequence) measuring magnetic field inhomogeneities before the functional scan. Images derived from this scan were used for registration between functional and anatomical images (Inati et al., 2011).

Image preprocessing

Functional images were preprocessed and analyzed using the SPM8 software package (Wellcome Department of Cognitive Neurology, London, UK). Because fourteen of the twenty subjects used for analysis moved their heads 0.5 mm or more during a single TR at least once during functional image acquisition (median number of large motions = 4.5), we used the Artifact Repair toolbox for SPM to correct for motion-related artifacts (Mazaika et al., 2009). First, runs of functional images were realigned using the standard SPM preprocessing pathway. Next, the time-series of motion-related distortions was estimated for each voxel separately and removed from the images. The Artifact Repair software estimates the effect of motion on the MR signal by regressing a time course of activity against a trigonometric function of the translational head motion estimates produced by realignment, and the resulting residuals from the regression represent the signal not explained by motion (Grootoonk et al., 2000). These voxel-wise residuals are written as images and used for all subsequent preprocessing and analysis. The Artifact Repair toolbox also repairs by interpolation volumes marred by very large motions (>1.5 mm) and volumes where the mean across voxels deviates more from the global mean than plausibly caused by physiological activity (>1.5% increase).

After motion correction, functional and structural images were coregistered by using intermediate registration targets derived from the multiecho GRE scan (Inati et al., 2011). That scan allows the estimation of undistorted spin density, field (B_0), and decay rate maps. The spin density map is used as a target for registration of the structural, and the other maps are used to produce a synthetic T_2*-weighted image, which has similar distortions and dropouts as the functionals, as a target for functional image registration. As these two targets are themselves in registration, this accomplishes cross-modal registration between the functional and anatomical scans.

Functional images were then normalized into a common space using a 12-parameter affine transformation estimated by warping the subject's anatomical image to match a template (SPM8 "segment and normalize"). Normalized functional images were resampled into 2×2×2 mm voxels and smoothed using a 7 mm full-width half-maximum Gaussian smoothing kernel.

Functional image analysis

All fMRI analyses were conducted in SPM8 using general linear models (GLM) fit by weighted least-squares. All regressors except the intercept terms in the models described below were convolved with SPM8's canonical hemodynamic response function, and the data were high-pass filtered with a cutoff of 128 s before estimation. To model individual subjects a separate set of regression coefficients was estimated for each run, and coefficients modeling the same effects in different runs were then summed to generate single-subject effect estimates. The subject-level effect estimates were then subjected to voxel-wise t-tests to produce *t*-maps for population-level inference. To determine statistical significance corrected for multiple comparisons, we threshold maps at $p < 0.005$ uncorrected and subject the resulting clusters of activated voxels to cluster-size correction controlling for the family-wise error rate (FWE) (Friston et al., 1993). Small volume cluster-size corrections based on anatomical masks were used when we had a priori hypotheses about the anatomical location of effects, and whole-brain correction when we did not. Reported p-values of fMRI activations reflect the corrected p-values of the clusters post-thresholding, rather than the uncorrected threshold used to generate maps.

Anatomical masks used for small-volume correction and generating anatomically-restricted functional ROIs were defined using the anatomical automatic labeling (AAL) atlas and generated via the WFU PickAtlas toolbox (Maldjian et al., 2003; Tzourio-Mazoyer et al., 2002), with the exception of one region for which we used a hand-drawn mask from a previous study. Specifically, for motor areas along the medial wall we used a mask consisting of the AAL's bilateral supplementary motor area and mid-cingulate cortex regions (8662 $2\times2\times2$ mm voxels); for the nucleus accumbens (which is not separately represented in AAL), a hand-drawn mask of the bilateral nucleus accumbens used originally in Daw (2011) following the guidelines of Breiter et al. (1997) (2374 voxels); and for vmPFC, AAL's anterior cingulate, medial orbital-frontal, and superior medial-frontal cortex regions (9414 voxels). Functional ROI analyses were performed by averaging estimated regression coefficients across all voxels within an ROI separately for each subject, then performing t-tests across subjects over the resulting averages. All activation maps in figures below, unless otherwise noted, are displayed at $p<0.005$ uncorrected.

Univariate GLMs

We constructed regressors modeling choice- and feedback-related activity using impulse functions at the shape onset and feedback onset times of successfully completed trials. To investigate neural correlates of valuation and reward-driven learning, we used the regressors V and δ as parametric modulators of the choice-time and feedback-time events respectively. V represents our estimate of the subject's value of the shape chosen on that trial, while δ is the difference between the actual reward received on that trial and that estimated value. These regressors were generated via simulation from the version of the RL model that we found to best describe the population, using a single set of parameters fit to all subjects simultaneously. A single representative set of parameters were used because we have repeatedly observed, in this and other data sets (Gershman et al., 2009; Simon and Daw, 2011b), that neural regressors generated using separate maximum-likelihood estimates of the parameters produce poorer fMRI results (i.e., noisier neural effect size estimates and diminished sensitivity). This is likely because parameters are not always well identified at the individual level, and variability in the point estimates effectively results in noisy rescaling of regressors between subjects, which in turn suppresses population level significance in fMRI (for additional discussion, see (Daw, 2011)). We also included as a parametric regressor of no interest at choice time the trial-by-trial estimates of choice history effects (C from Eq. (3)). Since subjects more often chose shapes associated with recent rewards, reward-based values are often correlated with choice history effects, so by including the choice history regressors we hoped to disentangle valuation-related activity from activity changes that might be associated with choosing the same options repeatedly. Finally, we also included a regressor modeling shape onsets for missed trials.

We used this GLM to identify regions of interest showing choice- and feedback-related activity without distinguishing between trials with different motor actions. This allowed us to subsequently make unbiased comparisons between value-related effects during bimanual, right- and left-handed actions within the selected ROIs. In order to compare value- and RPE-related activity during different motor actions, we fit a second GLM in which we divided single regressors from the above GLM into multiple regressors. To distinguish between different motor actions, we divided the choice-time event regressor into three different regressors modeling different types of trials: M_B for trials when the participant made a bimanual action, M_R for unimanual actions taken with the right hand, and M_L for unimanual actions with the left hand. We did the same with the feedback-time event regressor, splitting it into the three regressors F_B, F_R, and F_L. Similarly, to examine differences in value and RPE signals across the different types of trials we decomposed the chosen value regressor V into V_B, V_R, and V_L modulating M_B, M_R, and M_L respectively, and the δ regressor into δ_B, δ_R, and δ_L modulating F_B, F_R, and F_L respectively. As in the first GLM, we included a regressor for missed trials and regressors for choice history effects (also divided into separate regressors for bimanual, left, and right-handed choice trials) as regressors of no interest. (Note that the first GLM is equivalent to, but simpler to implement than, computing overall value effects by averaging together in a contrast the value estimates for each of the three motor action types from second GLM. The complexity of the latter approach is that the appropriate weighting for pooling the three trial types' parameter estimates—bimanual, left and right—depends, for each subject, on the particular number of times they chose left versus right.)

Finally, in addition to examining differences in chosen value signals across different motor movements, we also looked specifically for effector-specificity using right- and left-hand action values as defined in Palminteri et al. (2009), Wunderlich et al. (2009), and Hare et al. (2011). This model was identical to the second GLM described above, except instead of having chosen values V_R and V_L modulating activity on different trials, every trial in the unimanual condition (i.e. both event regressors M_R and M_L) was modulated by two regressors, one representing the values of the shapes that could be chosen by the right hand and the other representing the values of the shapes that could be chosen by the left hand. We combined the values of the two shapes choosable by a hand on any given trial into a single 'action value' regressor by simply averaging them. We used this model to compare only right- and left- handed unimanual actions, but not the unimanual and bimanual conditions, because there was no equivalent way to define action values in the bimanual case. This is because in the bimanual case it is not possible to segregate shape values into effector-specific action values because all actions are bimanual actions. Note that because bimanual and unimanual trials took place in separate sessions, each of which was modeled by distinct sets of regressors for each subject, not modeling the bimanual condition did not influence the estimation of the unimanual condition parameters.

Psycho-physiological interactions

To test for changes in patterns of functional connectivity between unimanual and bimanual choices, we fit two psycho-physiological interaction (PPI) models (Friston et al., 1997). The seed region for the PPI analyses was defined from a population-level $V>0$ contrast derived from the first GLM described above. We extracted a representative time course of activation from the seed region using SPM8, which takes the first eigenvector from a principle component analysis of all voxels within the seed region. That time course was then deconvolved against the HRF to extract the underlying neural activity, interacted multiplicatively with the (pre-convolution) M regressors modeling the choice-related activity at the time of shape onset, and then convolved again with the HRF. The resulting interaction regressors modeled activity related to the influence of the seed region specifically at the time of choice. Our first PPI model included two interaction terms: $S\times M_B$ for seed region influence during bimanual actions, and $S\times M_U$ for seed region influence during all unimanual (combined left and right handed) actions. Alongside the interaction terms we also included as nuisance variables the original seed region time series S and the entire set of regressors used in the second GLM described in the section above. The inclusion of every task-related regressor in the PPI model guards against observing spurious correlations between regions due to shared task-related effects rather than functional connectivity (O'Reilly et al., 2012). We then performed a whole-brain GLM, following the same methodology described above, and making population-level inferences using univariate t-tests over the contrasts of interest.

To ensure that the results of our first PPI were not being driven by unmodeled differences between seed region connectivity in right and left-handed actions, we performed a second PPI in which we included three interaction terms, $S\times M_R$, $S\times M_L$, and $S\times M_B$, modeling seed region influence during each of right, left, and bimanual actions. The model was otherwise identical to the first PPI model described above. For population-level inference, we entered estimates of the three $S\times M$

interaction terms into a factorial analysis for each voxel. This analysis modeled two factors: motor action type (right, left, or bimanual) and subject. Non-sphericity due to repeated measurements from the same subjects and differing amounts of data for different subjects and conditions was estimated and used to correct the final statistics (Penny and Henson, 2006). We then examined the conjunction of the $S \times M_B > S \times M_R$ & $S \times M_B > S \times M_L$ contrasts in the supplementary and cingulate motor areas along the medial wall of the brain, assessing significance using the conjunction null (Nichols et al., 2005). Significance was assessed using small-volume multiple comparison corrections, the small volume in question being the right and left hemisphere supplementary motor areas and mid-cingulate regions as defined by the Automatic Anatomical Labeling (AAL) atlas.

Results

Behavior

Participants earned an average of $10.81 (sd = $1) and failed to choose a shape in the allotted time on average 8.35 (sd = 5.83) out of 240 trials. We examined the behavioral data for signs that bimanual motions were more difficult than unimanual motions. There was no significant difference in the number of missed trials between conditions (difference mean = 0.65 trials, paired t-test, $t = -0.51, p > 0.6$). Average reaction time (RT) across participants did differ between motor action types, but rather than finding a consistent difference between unimanual and bimanual actions, we found that RTs on trials with right-handed actions (mean = 720 ms, sd = 62 ms) were faster than those with both left-handed (left hand RT mean = 763 ms, sd = 52 ms; paired samples difference, $t_{(19)} = 3.6$, $p < 0.01$) and bimanual (bimanual mean = 760 ms, sd = 73 ms; paired samples difference, $t_{(19)} = 2.8$, $p < 0.02$) actions, with no significant difference found between left-handed and bimanual actions ($t_{(19)} = 0.3$, $p > 0.75$).

Reinforcement learning model fitting

Since this study adds multi-effector response requirements to a four-bandit reinforcement learning task that has been well studied previously (Daw et al., 2006; Schönberg et al., 2007; Wimmer et al., 2012), we first verified that the learning behavior itself was well characterized by the model used in (Schönberg et al., 2007) and not substantially affected by the response manipulation. To compare how well models described the population we used the Bayesian random effects model selection procedure proposed in (Stephan et al., 2009) (implemented in spm_BMS from SPM8) to estimate the probability that an individual drawn from the population would be best described by a given model from a set. First, we validated the inclusion of both the reward learning parameters and the choice history parameters in our model by comparing the fit to choices of the full model to a model without an RL component (β_Q in Eq. (3) set to 0) and a model without a choice history component (β_C in Eq. (3) set to 0). The prevalence in the population of the model using both RL learning and choice history was estimated at 0.78, with an exceedance probability (the estimated chance of being the most common of the three models in the population) > 0.99. The model using only RL was estimated as the second most prevalent in the population with a probability of 0.17, while the model using only choice history without reward learning had an estimated prevalence of only 0.05. These results indicated that as in previous studies the model featuring both reward learning and choice history effects was best representative of the population (Lau and Glimcher, 2005; Schönberg et al., 2007).

Having established that the model featuring reward learning as well as choice history effects best described the population, we next tested for differences in the learning between conditions by comparing the fit of the RL model to the action-dependent RL model, which uses a separate set of parameters for trials in each condition. This analysis favored the unified model, with its estimated prevalence in the population at 0.95, and an exceedance probability > 0.99. These data suggest that the unified model provided the best description of the population. Accordingly—and in order to assure that any between-condition differences seen in the neuroimaging data were not artifacts of different model parameters used between blocks—we used the unified model with both choice-history and reward learning components as the basis for the fMRI analysis. The distribution across subjects of the unified RL model maximum-likelihood parameters, as well as the parameter values used to generate the fMRI regressors, are shown in Table 1.

Lateralization and coordination in the frontal cortex

In order to investigate the relationship between RL-related neural signals and multi-effector response requirements, we first examined the well-known lateralized organization of movement-related activity (Leyton and Sherrington, 1917; Penfield and Jasper, 1954; Sperry, 1961). As expected, greater activation during contralateral than ipsilateral hand movements (Fig. 2A) was found bilaterally in the precentral gyrus and central sulcus, extending backwards across the postcentral sulcus into the superior parietal lobe ($p < 0.001$; unless otherwise noted, fMRI statistics are quoted with cluster-size correction for family-wise error due to whole-brain multiple comparisons). Lateralized motor activity also appeared medial to the central sulcus, stretching across parts of supplementary motor area (SMA) and mid-cingulate cortex (MCC) ($p < 0.005$) (see Table 2 for a full list of significant activations).

While primary motor cortex is associated with effector-specific (lateralized) motor control, evidence from neuropsychology and neuroimaging suggest that motor areas along the medial wall such as the supplementary motor area are recruited in coordinating actions across the hemispheres (Brinkman, 1984; Jäncke et al., 2000; Laplane et al., 1977; Sadato et al., 1997). Accordingly, we examined movement-related activity in voxels along the medial wall of the frontal cortex that showed significant activity during the execution of choices regardless of the motor action used. As reported above, movement-related activity is lateralized even here: we found a significant hemisphere (right vs left) by hand (right vs left) interaction between the right and left hemisphere ROIs (paired samples, $t_{(19)} = 4.6, p < 0.001$). In addition to this lateralized pattern, we found evidence for augmented activity during bimanual actions. In particular, the right hemisphere ROI showed greater activation during bimanual than both right-handed (one-tailed test, $t_{(19)} = 3.2$, $p < 0.005$) and left-handed actions (one-tailed test, $t_{(19)} = 2.0, p < 0.05$). The left hemisphere ROI activation during bimanual actions was greater than during left-handed actions (one-tailed test, $t_{(19)} = 2.7, p < 0.01$), though it showed a weaker trend for being higher during right-handed actions (one-tailed test, $t_{(19)} = 1.3, p = 0.1$) (Fig. 2B). This pattern was different from activity in more lateralized motor areas (data not shown), where activity for bimanual action was similar to that for the contralateral movement alone. Thus, although regions along the medial wall share with the rest of the motor systems a lateralized structure, these regions may also be disproportionately recruited for coordinated action.

Value representations

We next searched for value representations complementary to the lateralized organization of the motor system. Lateralized value signals

Table 1
RL model parameter estimates.

	Individual fits—mean (s.e.m.)	Fixed effects fit—for fMRI
α	0.54 (0.06)	0.56
β_Q	2.94 (0.44)	1.8
κ	0.52 (0.07)	0.48
β_C	1.16 (0.3)	1.69

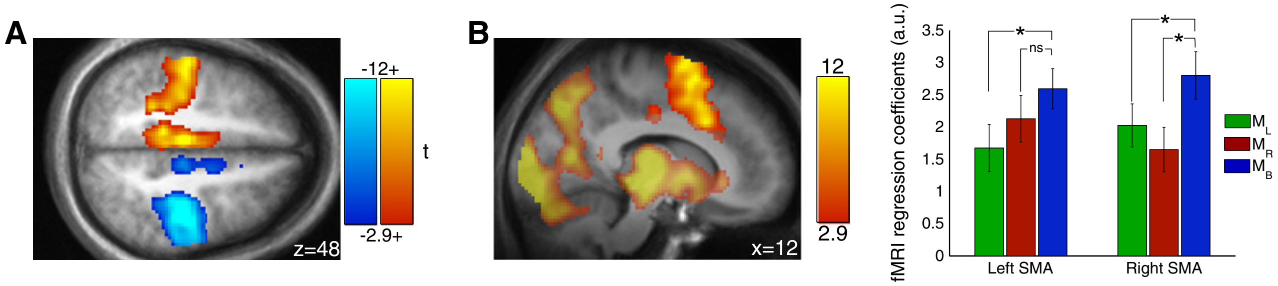

Fig. 2. Motor-related activity. (A) Regions of the bilateral precentral gyrus and medial motor cortex showing an effect of right versus left-handed actions. (B) The medial motor cortex ROI (SMA and MCC voxels showing activity across all motor conditions; see Methods section) showed greater activation during bimanual than either type of unimanual movement in the right hemisphere, with a similar trend in the left hemisphere.

would be appropriate for driving single-effector actions (Gershman et al., 2009; Palminteri et al., 2009), whereas value signals agnostic to effector or specific to the coordinated case could potentially be used to guide coordinated, multi-effector actions. Our approach was first to search for activity related to the value of the chosen option regardless of movement type, then examine its effector specificity (see Methods section). Note that testing whether activity in different regions is tied to different effectors requires comparing activity across regions (in our case, testing for an interaction of left or right hemisphere by left or right hand), necessitating the ROI approach.

We found value correlations in a region of strong prior hypothesis, the vmPFC ($p<0.05$, small-volume cluster size corrected within a medial prefrontal mask), and in a region of the medial frontal cortex extending from the SMA to the MCC ($p<0.001$ whole brain) (Fig. 3A). We also found activations unilaterally in the left hippocampus ($p<0.005$), the right hemisphere's lateral sylvian fissure extending ventrally from the rolandic operculum into the superior temporal sulcus ($p<0.05$), and the lateral prefrontal cortex dorsal to the vmPFC activation ($p<0.05$) (see Table 2). Although correlates of value in the hippocampus have previously been reported only sporadically, they tend to arise more often in tasks in which value computations require more reasoning about the relations between actions and events, perhaps in this case due to the complex response requirements (see Wimmer et al. (2012) for a full discussion).

We assessed effector-specificity of the value representations in the vmPFC and medial motor cortex by performing ROI analyses in which value was decomposed into regressors specific to particular types of motor actions. Two types of value signals have been used to determine effector-specificity in the fMRI literature; (Gershman et al., 2009) used chosen values, while (Palminteri et al., 2009) and (Wunderlich et al., 2009) used “action” values which model the value of each effector's actions on every trial, not just the chosen trials. These differ, effectively, in whether they model BOLD activity related to pre-choice values of available options (action values) or post-choice expected value (chosen values); either sort of value signal, to the extent it is present, might (or might not) be expressed in effector-specific regions. To examine effector specificity of value correlates on unimanual trials we used both types of value (see Methods section). To compare bimanual and unimanual choices (and for the remaining analyses in the paper, which largely concern the bimanual condition), we used only the chosen value regressors, since there was no comparable manner to define action values in the bimanual condition.

Examining the structure of the SMA/MCC activity, we identified two pairs of peaks aligned across the interhemispheric fissure. Each pair consisted of one peak in each hemisphere, with both peaks in a pair closely aligned with each other along the medial/lateral axis (see Table 2 for coordinates of the four peaks). Following a procedure analogous to that of Gershman et al. (2009), we split the ROI into lateralized pairs—anterior and posterior right and left hemispheres—and tested for hemisphere × hand interactions separately in the anterior and posterior ROI pairs. We found no lateralization in the posterior ROIs (paired samples, $t_{(19)}=-1.4$, $p>0.1$ Bonferroni corrected for multiple comparisons). The anterior ROI pair, however, did show value lateralization using both the action values ($t_{(19)}=3.8, p<0.005$, Bonferroni corrected for multiple comparisons) and chosen values ($t_{19}=2.5, p<0.05$, Bonferroni corrected for multiple comparisons), indicating that this region of cortex represented the action values of the contralateral hand more strongly than those of the ipsilateral hand (Fig. 3A, action values shown). Such lateralization for value in regions that also show lateralization for movements is consistent with the notion that the motor system contains effector-specific action value representations overlaid upon movement representations.

In contrast, similar lateralization was not detected in the value correlates in the vmPFC. Thresholded at $p<0.005$ uncorrected, the vmPFC cluster contained no voxels in the left hemisphere, so in order to test for lateralization in the vmPFC, we relaxed the threshold to $p<0.01$ uncorrected. The resulting cluster in the prefrontal cortex extended dorsally and laterally outside of the medial PFC, so we restricted the final ROI to the medial superior frontal gyrus, the anterior cingulate, and the medial orbital frontal cortex as defined by the AAL atlas. We found no significant hemisphere x hand interaction in the strength of action value representations in the vmPFC ROI using both the action values ($t_{(19)}=0.4, p>0.5$) and chosen values ($t=0.8, p=0.46$) (Fig. 3B, action values shown). Although null results should be interpreted with caution, this is consistent with our previous data (Gershman et al., 2009) and with theories that the vmPFC represents value over outcomes or goods rather than having an action- or effector-specific structure.

To be certain that we were not missing effector-specific value activations outside of our ROIs, we also examined the contrasts $V_R>V_L$ and $V_L>V_R$ over the whole brain (using both the action value and chosen value definitions). At whole-brain corrected thresholds, we found no significant difference between right- and left-hand values anywhere in the brain. However, this should not be interpreted as strong evidence against the presence of other lateralized value areas; in addition to the usual difficulties in interpreting null results, our study is underpowered for detecting effector-specificity at the whole-brain level due to our emphasis on the bimanual condition.

Finally, we compared value related activity (for chosen values) in the bimanual condition to that from the unimanual conditions. This activity was consistently in the same range as chosen-value related activities in the unimanual conditions. In neither the medial motor cortex ROI (minimum t-statistic, $t_{19}=0.3, p>0.5$), the vmPFC ROI ($t_{19}=-1.2, p>0.25$), nor anywhere in the brain did we find value-related activity that was significantly stronger in the bimanual condition than in both unilateral conditions. This suggested that there are not value areas that are differentially recruited for the bimanual case in this task, and led us to examine whether differences between the conditions might instead be observed in functional connectivity.

Table 2
fMRI activations. Maps were generated at $p<0.005$ uncorrected, and all activations that survived cluster-size whole-brain or small-volume correction at $p<0.05$ are reported.

	Region	Cluster size (2^3 mm voxels)	Peak (MNI)	*T*
Right > Left choices				
	Precentral sulcus & supplementary motor area L	5626	−40 −18 66	10.54
	Cerebellum & occipital lobe R	12,397	20 −48 −20	10.38
	Precentral sulcus R	3995	36 −16 60	−11.57
	Supplementary motor area R	582	8 −8 56	−6.54
	Cerebellum L	761	−22 −52 −22	−11.83
	Rolandic operculum R	861	42 −18 18	−8.08
	Occipital lobe L	360	−16 −94 20	−6.94
Chosen value				
	Rolandic operculum & superior temporal gyrus R	456	54 2 10	5.66
	Supplementary motor area & mid-cingulate cortex posterior L	1332	−12 −14 50	5.15
	Supplementary motor area & mid-cingulate cortex posterior R		12 −16 50	4.13
	Supplementary motor area & mid-cingulate cortex anterior L		−4 −2 46	4.94
	Supplementary motor area & mid-cingulate cortex anterior R		14 2 56	5.03
	Lateral prefrontal cortex L	381	−10 54 34	4.32
	Medial prefrontal cortex R[a]	195	6 60 −12	4.41
	Hippocampus L	582	−40 −12 −18	5.07
Reward prediction error				
	Occipital L	2554	−36 −84 −10	10.71
	Occipital R	2343	24 −90 0	9.61
	Inferior parietal lobe L	1256	−52 −56 48	9.24
	Lateral orbitalfrontal cortex R	657	46 46 −14	7.63
	Lateral orbitalfrontal cortex L	510	−44 50 −10	6.58
	Middle frontal gyrus L	581	−46 16 50	5.88
	Striatum L[a]	188	−10 8 −6	9.3
	Striatum R[a]	420	10 10 −8	5.98
vmPFC Functional connectivity—Bimanual > Unimanual				
	Mid-cingulate cortex R[a]	337	12 14 42	6.46
	Middle frontal gyrus R	772	52 30 18	5.03
vmPFC Functional connectivity—Bimanual > Right & Bimanual > Left				
	Mid-cingulate cortex R[a]	259	12 16 40	4.77

[a] Indicates small-volume correction used.

Reward prediction errors

We also looked for evidence of effector-specificity in reward prediction error (RPE) signaling in the striatum. BOLD signal in ventral striatum was correlated with RPEs at the time of reward feedback (right hemisphere $p<0.05$, left hemisphere $p<0.05$ small-volume cluster corrected within a nucleus accumbens mask) (Fig. 3C) (see Table 2 for a complete list of significant activations). Breaking the striatal activity down by effector as in Gershman et al. (2009), RPEs in this region showed an interaction between effector and hemisphere ($t_{(19)}=3.8, p<0.005$) such that prediction errors from unimanual choices were expressed relatively more strongly in the contralateral striatum. The RPE effect-size estimates (Fig. 3C) did not show a full crossover interaction where each hemisphere represented contralateral RPEs more strongly than ipsilateral RPEs, perhaps because there appeared to be, in addition to the effector × hemisphere interaction, an overall bias towards RPEs from right-hand actions. However, collapsing across both hemispheres, the main effect of right- vs left-hand choice RPEs did not reach significance ($t_{(19)}=1.5, p=0.14$). On bimanual trials, RPE signals fell between unimanual right and left RPE parameter estimates in both hemispheres, with no significant differences from either right or left-hand RPE signals in either hemisphere (minimum t-statistic, $t_{19}=-0.7, p=0.49$). A whole-brain analysis likewise showed no regions exhibiting significant differences between unimanual and bimanual trial prediction error expression.

PPI analysis

We next examined whether the shift from single- to multi-effector responses was accompanied by a change in functional connectivity. Motivated by the literature suggesting that this region instead codes the values of goals or outcomes more abstractly, and our own inability to detect effector specificity in the vmPFC value correlates in this and a previous study (Gershman et al., 2009), we hypothesized that these representations might be used preferentially to guide the selection of multi-effector actions. In this case, patterns of functional connectivity between the vmPFC and other regions of cortex should differ between the unimanual and bimanual conditions. More specifically, we predicted that regions involved in bimanual motor control, notably the supplementary motor and mid-cingulate cortices along the dorsal medial wall of the frontal cortex, would show stronger connectivity with the vmPFC during bimanual choices than unimanual choices.

To test this hypothesis we fit a whole-brain psycho-physiological interaction model with the vmPFC value ROI as the seed region, and contrasted the strength of functional connectivity with vmPFC in bimanual choices, $S\times M_B$, with the connectivity strength in unimanual choices, $S\times M_U$. In a targeted analysis of the medial motor regions, we found vmPFC connectivity was higher during bimanual relative to unimanual choices in the right anterior MCC ($p<0.05$, small-volume cluster size corrected within a medial motor region mask) (Fig. 4A). Additionally, at whole brain correction thresholds, the left lateral prefrontal cortex showed the same change in connectivity with vmPFC ($p<0.001$) (see Table 2).

This analysis maximized our power to find differences between bimanual and unimanual actions by pooling all unimanual actions together into a single trial type for which we had as many trials as we had bimanual trials. However, in principle this analysis might also incorrectly identify voxels where vmPFC connectivity changes depending only on whether or not a particular hand was being used. For example, a voxel might exhibit equivalent connectivity during both right-hand and bimanual actions and less connectivity during left-handed actions. Such a voxel might be detected by the $S\times M_B > S\times M_U$ contrast due to pooling across right and left hand trials. To rule out that explanation for our finding we ran a second analysis testing more conservatively for the conjunction of higher vmPFC connectivity during bimanual actions than each of right and left hand actions separately. Again, we

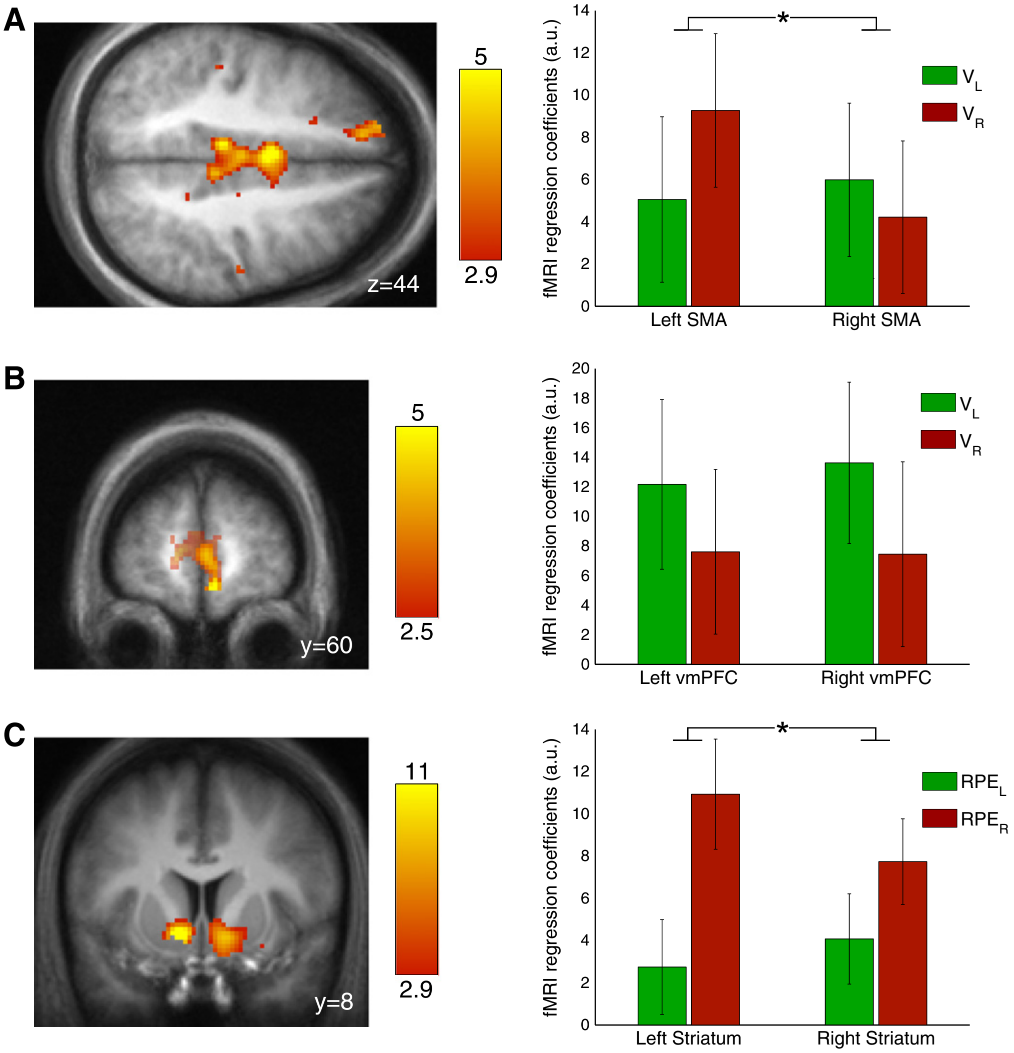

Fig. 3. Correlates of value and lateralization thereof. We found correlates of chosen value in the (A left) SMA and MCC and in the (B left) ventral–medial prefrontal cortex, as well as (C left) RPE signals in the ventral striatum. The medial motor activation contained two distinct pairs of contralateral peaks, so we divided the activation into four ROIs and tested for lateralization between contralateral ROIs. At our default threshold only the right vmPFC contained significant voxels, so in order to perform a similar test of value laterality we relaxed the threshold to $p<0.01$ (voxels active only at the lower threshold are displayed as semi-transparent). Examining both chosen value and action value effect sizes in these ROIs (action values shown), we find in the anterior SMA/MCC a significant hemisphere × effector interaction (A right), but not in the vmPFC (B right). We also find an effector × hemisphere interaction in the striatal RPE responses (C right).

found higher vmPFC connectivity during bimanual choices in the right anterior MCC ($p<0.05$, small volume cluster corrected within a medial motor region mask) (Fig. 4B). This analysis also revealed a trend towards condition-dependant connectivity in the left preSMA, though it did not reach significance (peak [0 -2 64], $T=3.53$; $p=0.07$, small-volume cluster corrected within a medial motor region mask; not reported in Table 2). However, under this analysis the lateral prefrontal cortex no longer showed significant connectivity changes.

Taken in conjunction with the other results, these motor-action dependant changes in functional connectivity suggest that communication between medial motor areas and the vmPFC increases during multi-effector actions. Although it is important to stress that a functional connectivity analysis of this sort does not speak to the direction or causality if the functionality connectivity, this change in connectivity patterns may facilitate the flexible mapping of decisions or decision variables from the vmPFC onto complex multi-effector actions.

A further question is whether the medial area identified in the PPI analysis is the same area identified above, for lateralized value. At the threshold we examined, we did not find these two effects significant in the same voxels. The ROI exhibiting lateralized value signals (Fig. 3A) lies approximately 16 mm posterior to the site of vmPFC connectivity, with no overlap between them. Of course a lack of overlapping voxels is not sufficient evidence to claim a dissociation between these two areas (Henson, 2005), and we therefore cannot draw strong conclusions as to whether or not the unimanual value and bimanual connectivity signals share precisely the same neural substrate in the medial motor regions. However, the site of bimanual vmPFC connectivity does fall within the mid-cingulate region showing greater overall bimanual than unimanual motor activity (Fig. 2B), suggesting a convergence of motor signals and motor-dependant connectivity in this area. More generally, the proximity of the motor, value, and connectivity results reported here suggests a convergence along the medial wall of motor functions and signals that could be used to guide them.

Discussion

In this study we examined how models of action evaluation by reinforcement learning in the brain, which have typically been formulated in the context of choices made using single effectors, scale up to richer, more realistic settings in which multiple effectors must act in concert to achieve a single goal. Beyond the problems of high dimensional movement control and coordination per se, reinforcement learning in this

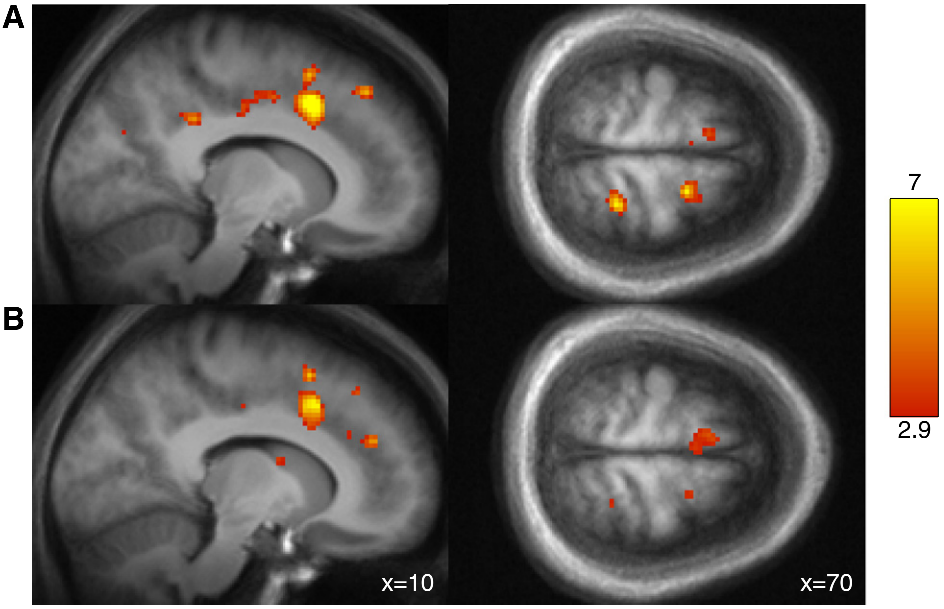

Fig. 4. Functional connectivity with the vmPFC. (A) Voxels in the SMA/MCC showing higher vmPFC connectivity during bimanual choices than unimanual choices overall, collapsing across right and left-handed choices. (B) The conjunction of bimanual connectivity greater than right and bimanual greater than left-hand action connectivity, verifying that the effects are not driven by asymmetries between right and left action connectivity strengths. While we find activity in both the MCC and SMA, only the MCC survives small-volume cluster correction.

setting requires assigning values to a large space of actions, and updating all these values from experience. This is a classic curse of dimensionality. We (Gershman et al., 2009) have previously examined a special case of this problem: multi-effector actions arising from responses that, although simultaneous, are rewarded independently, which would allow the brain evaluate and choose actions separately for each effector.

In the present study, we examine the opposing case where both hands must be used in an inseparable fashion by having subjects chose among options in a bandit task via bimanual button presses. This defeats the effector-specific mechanisms posited by Gershman et al. (2009) by making it impossible to evaluate the reward expected for a movement with one hand separate from the other hand's movement. We contrasted neural activity during this bimanual condition against a more traditional unimanual condition, in which each option corresponded to a single button press using one or the other hand. Consistent with previous reports that sensorimotor and motor-planning regions contain effector-specific values (Gershman et al., 2009; Platt and Glimcher, 1999; Wunderlich et al., 2009), we found that value signals in the medial motor cortex showed contralateral biases indicating effector specificity, whereas value signals in the vmPFC did not exhibit any such effector-specific laterality effects. We also found, as in Gershman et al. (2009), that reward prediction error signaling in the ventral striatum reflects more strongly choices executed by the contralateral hand, providing further evidence that the brain contains decision and reward learning systems that are subdivided by effectors. Though right medial motor cortex was overall more active during bimanual movements, we found no evidence there or elsewhere for neural representations of value that were unique or preferential for bimanual actions. However, bimanual and unimanual conditions were associated with different patterns of connectivity between another area with value correlates and motor control regions, with the vmPFC showing greater connectivity with the medial motor cortex during bimanual choices than during choices that could be executed by single effectors. This change in connectivity suggests that the brain may, in this task, have resolved the challenge of evaluating multi-effector actions by performing valuation in the space of abstract “goods” and using top-down connectivity with motor planning areas to convert these decisions into movements. It is important to keep in mind that this interpretation of the observed changes in functional connectivity remains speculative, as our functional connectivity measures do not assess whether information is being communicated from the vmPFC to the MCC, in the opposite direction, or indeed in both directions simultaneously or from a common afferent. Future experiments in multi-effector decision making using single-unit electrophysiology or imaging methods with higher temporal resolution, such as EEG or MEG, may help resolve this issue, as well as determine the nature of the information being communicated.

Of course we do not rule out the possibility that motor and motor control areas also contain substrates for multi-effector action value representations. Though in this task we did not find evidence for differential representation of multi-effector value at the level of BOLD signaling, such representations are plausible. For instance, the brain's motor systems do contain neurons tuned for particular multi-effector actions Tanji et al. (1988), Donchin et al. (2002), and such representations would also be a natural substrate or precursor for values over these actions as well. However, the curse of dimensionality makes this impractical as a general solution for all possible multi-effector actions, so we speculate that such coding would be most useful for small groups of effectors that are often coupled during normal behavior (such as hands and eyes), or for familiar, well-practiced multi-effector actions. Consistent with this proposal, several studies have indicated that representations of novel actions in motor cortex develop over the course of motor learning and practice, while prefrontal involvement decreases with practice (Franz et al., 2000; Jenkins et al., 1994; Karni et al., 1995; Shadmehr, 1997), suggesting a shift from top-down control for novel actions to explicit motor representations for familiar ones. Thus, decision making over very familiar or well-practiced bimanual actions in particular might plausibly involve action-specific, if not effector-specific, motor representations.

Separability and coordination in bimanual actions

The present study is far from the first in psychology and neuroscience to investigate bimanual action. However, the present study builds on work on decision-making to investigate valuation over multi-effector movements, and so it concerns rather different questions and phenomena than much of the literature on bimanual motor coordination. The classic and most commonly studied topic in bimanual action in psychology concerns overcoming biases that oppose and impair coordinated action. Specifically, the common finding is that asymmetric bimanual actions are, compared to unimanual actions, slower and more prone to error (Kelso et al., 1979). This has been interpreted as evidence that cross-talk, or possibly shared motor plans between the two hands, biases isomorphic muscle groups in both hands to be activated together in 'mirrored' actions (Kelso et al., 1979; Oliveira and Ivry, 2008; Swinnen and Wenderoth, 2004). We consider this issue to be distinct from that investigated here because impairments for bimanual actions are not universal, appearing only during particular circumstances and motor action types (Mechsner et al., 2001; Oliveira and Ivry, 2008; Rosenbaum et al., 2006; Swinnen and Wenderoth, 2004), and indeed we do not observe any such impairments in performance or reaction times in the current study. Bimanual penalties occur reliably when actions are self-generated or cued by arbitrary symbols, but simple manipulations like introducing visual cues or targets can eliminate such penalties, perhaps by allowing actions to be guided by ocular-motor pathways or encouraging unified representations of an action (Diedrichsen et al., 2006; Oliveira and Ivry, 2008). The current task may have not elicited bimanual penalties because the shapes acted as visual targets, because the choosing of a shape was represented as a "unified" motor action, or because the motor actions involved were sufficiently simple. In any case, the challenge motivating the current study is not at the level of generating hand-specific movements but instead, given the ability to do so, at the level of evaluation: how the brain evaluates and decides between options when action values cannot be separated into hand-specific terms.

Prefrontal value regions in unimanual and bimanual actions

By design, our experiment does not produce a true curse of dimensionality, since there are only four options in either condition. Indeed, this is key to our interpretation of our results, since we envision that the low dimensionality of the options is what enables the brain to evaluate the options, even though the movements required to execute a choice couple multiple effectors. It is worth considering several caveats to this interpretation.

First, in line with much work on vmPFC and OFC value correlates across species, we assume that these systems represent value more abstractly—perhaps relative to particular goals, goods, or options, such as shapes here—rather than over actions per se (Gläscher et al., 2009; Levy and Glimcher, 2011; Padoa-Schioppa and Assad, 2006; Rudebeck et al., 2008; Wallis and Miller, 2003; Wunderlich et al., 2010). Consistent with this (and with our previous effort, Gershman et al. (2009)), we do not detect a bias toward contralateral value representations in vmPFC, of the sort that we detect in more posterior cortices and associate with effector-specific valuation. However, as this is ultimately a null result, it is particularly important to consider two previous studies that do appear to detect some effector specificity in vmPFC.

Notably, Palminteri et al. (2009), using a task similar to the unimanual condition in the current study, reports value-signal lateralization in the ventral PFC and OFC. An important aspect of these data, however, is that laterality measures are pooled over areas of ventral PFC comprising both a medial activation similar to that reported here (and those that are the focus of most of the work in neuroeconomics discussed above), plus another set of distinct activations in the bilateral lateral OFC (and not detected in the current study). From the activation maps presented, it appears as though this more lateral region comprises the bulk of the ROI used and thus may be the primary source of effector specificity in the Palminteri et al. (2009) data, whereas the medial site, as in our hands (see also (Gershman et al., 2009)), may not show such an effect. Indeed, recent evidence from electrophysiology, lesion, and imaging studies suggest that the lateral OFC and vmPFC are functionally distinct subregions, with the former involved in associated stimuli and outcomes and the latter representing values of available options and driving decisions (Noonan et al., 2010; Wallis, 2012; Walton et al., 2010), so it is possible that these regions also differ in their effector-specificity.

Also consistent with this interpretation are the results of Hare et al. (2011), which also had participants choosing between two options by pressing a button with either their right or left hand. As in the current study, they found action value signals only in vmPFC and not in lateral OFC, and find no indication of effector-specificity in the vmPFC signals. Finally, Wunderlich et al. (2009) demonstrated adjacent, but somewhat disjoint, activation maps for value in vmPFC, depending whether the values were associated with a hand or eye movement. However, the temptation to interpret these activations as evidence for an effector-specific dissociation in this region does not appear to be supported statistically by a significant interaction of region by effector. If the apparent difference in effects between these subregions is not itself significant, then the data would in fact be consistent with the effector-nonspecific interpretation given here (Henson, 2005; Nieuwenhuis et al., 2011; Poldrack et al., 2008).

A second subtlety about the present task is that although we manipulate the response mode to require either one- or two-handed button presses, in both conditions value is tied to the shapes, whose button mappings permute from trial to trial, rather than to the hand movements themselves. Because of this, even in our unimanual condition, learning and evaluation cannot logically take place entirely at the level of effector-specific value maps, but must be grounded in a shape-value mapping to carry value from trial to trial. Under our interpretation, this mapping could implicate vmPFC in both conditions. However, a potentially key difference between the conditions is that in the unimanual condition—but not the bimanual one—learned action values could be projected onto an effector-specific value map, and decision ultimately conducted in that action value space (Glimcher, 2008). (As discussed further below, such a decision would involve competition across the hemispheres to determine whether the left or right hand responded, but, again, unlike in the bimanual case, these interactions could be grounded in effector-specific value maps.) This is consistent with work on saccadic decision making in primates (Sugrue and Corrado, 2004), in which the value of a saccade to a location is typically signaled by some feature of the target other than its location, such as color. Although neurons in posterior parietal area LIP correlate with the value of saccades to particular locations, these values likely are communicated from elsewhere in these tasks because LIP neurons are not generally tuned for the information about color that is required to determine the value of different actions. Notwithstanding its origin, though, this "value map" over saccade locations is widely viewed as a candidate causal contributor to saccadic decisions (Glimcher, 2008; Sugrue and Corrado, 2004). In the case of the present study, even if vmPFC contributes the shape-value mapping in both conditions, its top-down involvement may be relatively stronger in the bimanual condition, since in this case the decision itself also cannot be conducted within or across posterior effector-specific representations. A broader point is that many studies support the idea that the relative involvement of prefrontal (e.g., outcome-based) and more posterior (e.g., action-based) decision mechanisms differ depending on circumstances; what our results here suggest is that a multi-effector response requirement also impacts this balance. For instance, unlike in the present study, Hare et al. (2011) emphasize connectivity between the vmPFC and SMA in guiding even unimanual actions. Although this is not inconsistent with our finding of greater connectivity in a bimanual than unimanual conditions, increased involvement of top-down, "goal-directed" valuation might also arise from that study's use of choices between different types of juice conveyed symbolically rather than acquired by reinforcement learning, as in the current study. Indeed, value correlates in vmPFC are famously sensitive to subjective preferences that are grounded in very high-level, symbolically cued information such as branding (Harvey et al., 2010; McClure et al., 2004; Plassmann et al., 2008). Evidence from rats and humans indicates that goal-directed decision-making relies on value signals originating in the prefrontal cortex, while decisions guided by reinforcement learning may not (Camille et al., 2011; McDannald, et al., 2011; Rudebeck and Murray, 2011). Similarly, the prefrontal cortex is known to play a major role in decision-making that requires executive control or working memory capacity, regardless of the number of effectors used McClure et al. (2004), O'Reilly and Frank (2006), Frank et al. (2007), Hare et al. (2009), while imaging results from a reinforcement learning task indicated the presence of decision variables in the SMA and preSMA rather than vmPFC (Wunderlich et al., 2009).

A related set of studies showed that the roles of prefrontal cortex and medial motor cortex in single-effector choice also depend on whether rewards are associated with visual stimuli versus particular motor actions without any accompanying stimuli. BOLD correlates of value in SMA and mid-cingulate cortex are stronger during pure action choices than choices over visual stimuli (Gläscher et al., 2009). Similarly, lesions of dorsal ACC and anterior mid-cingulate cortex impaired reward learning from pure actions but not visual stimuli, whereas OFC lesions impaired only learning of stimulus values (Gläscher et al., 2009; Kennerley et al., 2006; Rudebeck et al., 2008). These studies emphasize the importance of medial motor regions in action valuation and selection, as well as suggest that a number of variables can modulate the relative contributions of motor and prefrontal areas during even unimanual action.

All these results are broadly consistent with our interpretation here, but they do underline the question why, if the bimanual condition implicates 'top-down' control differentially, we did not detect differential recruitment of vmPFC (or any other) value correlates in a univariate sense, but instead only changes in functional connectivity. One explanation for the lack of a difference in vmPFC value correlates in the bimanual and unimanual conditions is that while 'top-down' influences on value may require the vmPFC, the vmPFC is still capable of representing values in simpler conditions in which it may not be required. Indeed, a "goods-based" valuation system can represent everything an effector-specific action value representation can. The vmPFC may track value in parallel to other systems, even when the other systems are preferentially controlling actual choices, for instance in order to facilitate behavioral flexibility should a change in environment render such flexibility necessary and the vmPFC's representations more useful. (See Daw et al. (2005) and Simon and Daw (2011b) for further discussion of valuation systems working in parallel.)

Overall, the broader literature appears to be consistent with the suggestion, as here, that these systems support action- vs outcome-based decision making, and the current results thus are consistent with the idea that the latter should be preferentially engaged for multi-effector action. In the present study we have not directly modulated all the factors that would be necessary to directly test the involvement of action- or outcome-tied representations, per se, e.g. by manipulating the presence or absence of value-carrying stimuli, or testing goal-directedness via outcome devaluation (Dickinson and Balleine, 2002). These investigations are left for future work.

Medial motor regions in bimanual actions and action selection

The medial motor cortex is commonly associated with bimanual coordination (Brinkman, 1984; Jäncke et al., 2000; Laplane et al., 1977; Sadato et al., 1997). However, lesion studies suggest that in addition to their involvement in bimanual actions, the medial motor cortex mediates communication between hemispheres during unimanual actions as well (Brinkman, 1984; Laplane et al., 1977). Lesions to the SMA and preSMA often result in impairments in which the hand contralateral to the lesioned hemisphere mirrors the actions of the other. Thus, one reason these areas might be implicated in even the unimanual condition of our task—and also those of Palminteri et al. (2009) and Hare et al. (2011)—is that all these tasks involve a choice entered with one hand or the other. Thus, although each effector could determine the values of its actions independently, ultimately one hand must move and the other must not. Such competitive interactions clearly implicate communication across the hemispheres—presumably less than in our bimanual condition, in which both hands must move in concert, and perhaps more than in other bandit tasks (e.g., Daw et al. (2006)) in which only right hand responses were used. These considerations may also explain why in the present task we find effector-specific value correlates in medial motor regions (see also Wunderlich et al. (2009)), whereas in Gershman et al. (2009) we found similar effects in the parietal lobe. The difference may be that although Gershman et al. (2009) used only bimanual movements, these were built up by combining two independent unilateral decisions, and thus does not involve the same sort of competitive interactions as the unimanual task here. Having effector-specific values in the medial motor cortex may facilitate the communication necessary for these interactions.

Given the association between the medial motor cortex and bimanual action, which we replicate here, one possible interpretation of our results is that the vmPFC communicates with medial motor areas preferentially in the bimanual condition specifically because those areas are responsible for generating bimanual actions. However, recent imaging studies have emphasized that rather than being specialized for bimanual actions, these medial motor areas, particularly those more anterior, contribute generally to planning complex motor actions and higher cognitive functions. For instance, Hare et al. (2011)) view them as a comparator and decision region, and a relay between frontal valuation regions and primary motor cortex, while Kolling et al. (2012) and Hayden et al. (2011)) find them implicated in monitoring rewards and signaling exploration during foraging tasks. The preSMA in particular is thought to be involved in relatively abstract motor planning and planning in general, such as in preparing sequences of actions and encoding stimulus-action associations (Picard and Strick, 2001; Simon and Daw, 2011a), while the mid-cingulate cortex has a wide array of cognitive and motor functions which we survey briefly below. One recent study of bimanual reaching (Diedrichsen et al., 2006) found that the preSMA and anterior MCC showed increased activation not during bimanual actions generally, but only non-symmetric actions that were not spatially guided. This led the authors to propose that the role of these areas in bimanual motor control is an extension of their roles in motor control and cognitive control generally. Thus, the vmPFC-preSMA/aMCC connectivity observed here may not be specific to bimanual choices. Coupling between value and motor-planning regions may instead be common to any decision making situation in which the motor actions involved are complex enough to require preSMA and aMCC engagement.

The region of medial motor cortex where we find strongest action-dependent connectivity to frontal value regions is the mid-cingulate cortex. Cingulate cortex contains a number of anatomical and functional subdivisions. The area identified here is commonly referred to as rostral cingulate zone (RCZ) or anterior mid-cingulate cortex (aMCC) in the motor literature (Picard and Strick, 2001; Shackman et al., 2011), and sometimes the dorsal anterior cingulate in the decision-making literature (Kolling et al., 2012). Anatomical and functional connectivity analyses indicate that this region is connected to lateral and medial prefrontal cortex, subcortical nuclei including the striatum and amygdala, and primary and premotor cortices (Beckmann et al., 2009; Yu et al., 2011). Functional MRI studies have implicated this region in a number of different processes, most prominently cognitive control, processing of aversive, emotional, and rewarding stimuli, and motor action selection (Büchel et al., 2002; Martin and Delgado, 2011; Phelps et al., 2004; Seo and Lee, 2007; Shackman et al., 2011; Venkatraman et al., 2009; Xue et al., 2008). This has led to the hypothesis that this region serves as a hub linking value and affects signals from prefrontal and subcortical regions to cognitive and motor control systems (Shackman et al., 2011). Consistent with this view, during decision making tasks, neurons in the primate homologue of the aMCC carry information both about decision variables and motor actions, though reports differ as to whether motor and value signals coexist in single neurons or in distinct populations (Hayden and Platt, 2010; Kennerley and Wallis, 2009; Luk and Wallis, 2009; Matsumoto et al., 2003). The aMCC thus seems ideally situated functionally and anatomically to act as a conduit through which decision variables can exert control over the motor system.

Conclusions

Together with previous studies of value correlates and connectivity during decision-making, our results suggest that the roles of different cortical areas in guiding actions depend on the details of the task at hand. When valuation depends on current goals and internal states, as in goal-directed choice, or cannot be represented easily by effector-specific actions, as in bimanual coordination, frontal representations appear to be critical for guiding choice. When values can be mapped directly onto stimulus-action associations, value representations in posterior visuo-motor and motor control areas may play a more prominent role in action selection. Adaptively subdividing decision problems appropriate to their structure may help to ameliorate the curse of dimensionality problems that arise for decisions over multiple effectors. However, if the brain indeed adapts its approach and systems used to task circumstances, a further question is what are the mechanisms and principles by which it does so. Although there exist proposals for how the brain might arbitrate between distinct decision-making systems with overlapping domains (Daw et al., 2005; Dickinson, 1985; Poldrack et al., 2001), the general problem of adaptive dimensionality reduction for reinforcement learning remains largely open.

Acknowledgments

This work was supported by National Institute of Mental Health Grant R01 MH07882 as part of the National Science Foundation–National Institutes of Health Collaborative Research in Computational Neuroscience Program, by Human Frontiers Science Program Grant RGP0036/2009-C, a Scholar Award from The McKnight Foundation and a Scholar Award from The McDonnell Foundation. The authors would like to thank Dylan Alex Simon for helpful comments and discussion.

References

Alexander, G.E., DeLong, M.R., Strick, P.L., 1986. Parallel organization of functionally segregated circuits linking basal ganglia and cortex. Annu. Rev. Neurosci. 9, 357–381.

Andersen, R.A., Buneo, C.A., 2002. Intentional maps in posterior parietal cortex. Annu. Rev. Neurosci. 25, 189–220.

Batista, A.P., Buneo, C.A., Snyder, L.H., Andersen, R.A., 1999. Reach plans in eye-centered coordinates. Science 285 (5425), 257–260.

Beckmann, M., Johansen-Berg, H., Rushworth, M.F.S., 2009. Connectivity-based parcellation of human cingulate cortex and its relation to functional specialization. J. Neurosci. 29 (4), 1175–1190.

Berns, G.S., McClure, S.M., Pagnoni, G., Montague, P.R., 2001. Predictability modulates human brain response to reward. J. Neurosci. 21 (8), 2793–2798.

Blanke, O., Spinelli, L., Thut, G., Michel, C.M., Perrig, S., Landis, T., Seeck, M., 2000. Location of the human frontal eye field as defined by electrical cortical stimulation: anatomical, functional and electrophysiological characteristics. Neuroreport 11 (9), 1907–1913.

Breiter, H.C., Gollub, R.L., Weisskoff, R.M., Kennedy, D.N., Makris, N., Berke, J.D., Goodman, J.M., Kantor, H.L., Gastfriend, D.R., Riorden, J.P., Mathew, R.T., Rosen, B.R., Hyman, S.E., 1997. Acute effects of cocaine on human brain activity and emotion. Neuron 19 (3), 591–611 (Sep).

Brinkman, C., 1984. Supplementary motor area of the monkey's cerebral cortex: short- and long-term deficits after unilateral ablation and the effects of subsequent callosal section. J. Neurosci. 4 (4), 918–929.

Büchel, C., Bornhovd, K., Quante, M., Glauche, V., Bromm, B., Weiller, C., 2002. Dissociable neural responses related to pain intensity, stimulus intensity, and stimulus awareness within the anterior cingulate cortex: a parametric single-trial laser functional magnetic resonance imaging study. J. Neurosci. 22 (3), 970–976.

Calton, J.L., Dickinson, A.R., Snyder, L.H., 2002. Non-spatial, motor-specific activation in posterior parietal cortex. Nat. Neurosci. 5 (6), 580–588.

Camille, N., Griffiths, C.A., Vo, K., Fellows, L.K., Kable, J.W., 2011. Ventromedial frontal lobe damage disrupts value maximization in humans. J. Neurosci. 31 (20), 7527–7532.

Daw, N.D., 2011. Trial-by-trial data analysis using computational models. In: Delgado, M.R., Phelps, E.A., Robbins, T.W. (Eds.), Affect, Learning and Decision Making, Attention and Performance XXIII. Oxford University Press, Oxford.

Daw, N.D., Niv, Y., Dayan, P., 2005. Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. Nat. Neurosci. 8 (12), 1704–1711.

Daw, N.D., O'Doherty, J.P., Dayan, P., Seymour, B., Dolan, R.J., 2006. Cortical substrates for exploratory decisions in humans. Nature 441 (7095), 876–879.

Deichmann, R., Gottfried, J., Hutton, C., Turner, R., 2003. Optimized EPI for fMRI studies of the orbitofrontal cortex. Neuroimage 19 (2), 430–441.

Dickinson, A., 1985. Actions and habits: The development of behavioural autonomy. Philos. Trans. R. Soc. Lond. B Biol. Sci. 308 (1135), 67–78.

Dickinson, A., Balleine, B., 2002. The role of learning in the operation of motivational systems, In: Gallistel, C.R. (Ed.), Stevens Handbook of Experimental Psychology: Learning, Motivation and Emotion, 3rd edition. John Wiley & Sons, New York, pp. 497–534.

Diedrichsen, J., Grafton, S., Albert, N., Hazeltine, E., Ivry, R.B., 2006. Goal-selection and movement-related conflict during bimanual reaching movements. Cereb. Cortex 16 (12), 1729–1738.

Donchin, O., Gribova, A., Steinberg, O., Mitz, A.R., Bergman, H., Vaadia, E., 2002. Single-unit activity related to bimanual arm movements in the primary and supplementary motor cortices. J. Neurophysiol. 88 (6), 3498–3517.

Frank, M.J., Scheres, A., Sherman, S.J., 2007. Understanding decision-making deficits in neurological conditions: insights from models of natural action selection. Philos. Trans. R. Soc. Lond. B Biol. Sci. 362 (1485), 1641–1654.

Franz, E.A., Waldie, K.E., Smith, M.J., 2000. The effect of callosotomy on novel versus familiar bimanual actions: a neural dissociation between controlled and automatic processes? Psychol. Sci. 11 (1), 82–86.

Friston, K.J., Worsley, K.J., Frackowiak, R.S.J., Mazziotta, J.C., Evans, A.C., 1993. Assessing the significance of focal activations using their spatial extent. Hum. Brain Mapp. 1 (3), 210–220.

Friston, K.J., Buechel, C., Fink, G.R., Morris, J., Rolls, E., Dolan, R.J., 1997. Psychophysiological and modulatory interactions in neuroimaging. Neuroimage 6 (3), 218–229.

Gershman, S.J., Pesaran, B., Daw, N.D., 2009. Human reinforcement learning subdivides structured action spaces by learning effector-specific values. J. Neurosci. 29 (43), 13524–13531.

Gläscher, J., Hampton, A.N., O'Doherty, J.P., 2009. Determining a role for ventromedial prefrontal cortex in encoding action-based value signals during reward-related decision making. Cereb. Cortex 19 (2), 483–495.

Glimcher, P.W., 2008. Choice: towards a standard back-pocket model. In: Glimcher, P.W., Camerer, C.F., Fehr, E., Poldrack, R.A. (Eds.), Neuroeconomics: Decision Making and the Brain. Academic Press, San Diego, CA, pp. 503–521 (Ch. 31).

Grootoonk, S., Hutton, C., Ashburner, J., Howseman, A.M., Josephs, O., Rees, G., Friston, K.J., Turner, R., 2000. Characterization and correction of interpolation effects in the realignment of fMRI time series. Neuroimage 11 (1), 49–57.

Hare, T.A., Camerer, C.F., Rangel, A., 2009. Self-control in decision-making involves modulation of the vmPFC valuation system. Science 324 (5927), 646–648.

Hare, T.A., Schultz, W., Camerer, C.F., O'Doherty, J.P., Rangel, A., 2011. Transformation of stimulus value signals into motor commands during simple choice. Proc. Natl. Acad. Sci. U.S.A. 108 (44), 18120–18125.

Harvey, A.H., Kirk, U., Denfield, G.H., Montague, P.R., 2010. Monetary favors and their influence on neural responses and revealed preference. J. Neurosci. 30 (28), 9597–9602 (Jul).

Hayden, B.Y., Platt, M.L., 2010. Neurons in anterior cingulate cortex multiplex information about reward and action. J. Neurosci. 30 (9), 3339–3346.

Hayden, B.Y., Pearson, J.M., Platt, M.L., 2011. Neuronal basis of sequential foraging decisions in a patchy environment. Nat. Neurosci. 14 (7), 933–939.

Henson, R., 2005. What can functional neuroimaging tell the experimental psychologist? Q. J. Exp. Psychol. 58A (2), 193–233.

Inati, S., Velasco, P., Vessel, E., 2011. A new method for robust functional to structural alignment using multi-echo GRE B0 and R2* mapping. Poster at Hum. Brain Mapp. (Quebec City (Canada)).

Jäncke, L., Peters, M., Himmelbach, M., Nösselt, T., Shah, J., Steinmetz, H., 2000. fMRI study of bimanual coordination. Neuropsychologia 38 (2), 164–174.

Jenkins, I., Brooks, D., Nixon, P., Frackowiak, R., Passingham, R., 1994. Motor sequence learning: a study with positron emission tomography. J. Neurosci. 14 (6), 3775–3790.

Kable, J.W., Glimcher, P.W., 2007. The neural correlates of subjective value during intertemporal choice. Nat. Neurosci. 10 (12), 1625–1633.

Karni, A., Meyer, G., Jezzard, P., Adams, M.M., Turner, R., Ungerleider, L.G., 1995. Functional MRI evidence for adult motor cortex plasticity during motor skill learning. Nature 377 (6545), 155–158.

Kelso, J.A., Southard, D.L., Goodman, D., 1979. On the coordination of two-handed movements. J. Exp. Psychol. Hum. Percept. Perform. 5 (2), 229–238.

Kennerley, S.W., Wallis, J.D., 2009. Encoding of reward and space during a working memory task in the orbitofrontal cortex and anterior cingulate sulcus. J. Neurophysiol. 102 (6), 3352–3364.

Kennerley, S.W., Diedrichsen, J., Hazeltine, E., Semjen, A., Ivry, R.B., 2002. Callosotomy patients exhibit temporal uncoupling during continuous bimanual movements. Nat. Neurosci. 5 (4), 376–381.

Kennerley, S.W., Walton, M.E., Behrens, T.E.J., Buckley, M.J., Rushworth, M.F.S., 2006. Optimal decision making and the anterior cingulate cortex. Nat. Neurosci. 9 (7), 940–947.

Knutson, B., Westdorp, A., Kaiser, E., Hommer, D., 2000. FMRI visualization of brain activity during a monetary incentive delay task. Neuroimage 12 (1), 20–27.

Kolling, N., Behrens, T.E.J., Mars, R.B., Rushworth, M.F.S., 2012. Neural mechanisms of foraging. Science 336 (6077), 95–98.

Laplane, D., Talairach, J., Meininger, V., Bancaud, J., Orgogozo, J., 1977. Clinical consequences of corticectomies involving the supplementary motor area in man. J. Neurol. Sci. 34 (3), 301–314.

Lau, B., Glimcher, P.W., 2005. Dynamic response-by-response models of matching behavior in rhesus monkeys. J. Exp. Anal. Behav. 84 (3), 555–579.

Levy, D.J., Glimcher, P.W., 2011. Comparing apples and oranges: using reward-specific and reward-general subjective value representation in the brain. J. Neurosci. 31 (41), 14693–14707.

Levy, I., Schluppeck, D., Heeger, D.J., Glimcher, P.W., 2007. Specificity of human cortical areas for reaches and saccades. J. Neurosci. 27 (17), 4687–4696.

Leyton, A.S.F., Sherrington, C.S., 1917. Observations on the excitable cortex of the chimpanzee, orangutan, and gorilla. Q. J. Exp. Physiol. 11 (2), 135–222.

Luk, C.-H., Wallis, J.D., 2009. Dynamic encoding of responses and outcomes by neurons in medial prefrontal cortex. J. Neurosci. 29 (23), 7526–7539.

Maldjian, J.A., Laurienti, P.J., Kraft, R.A., Burdette, J.H., 2003. An automated method for neuroanatomic and cytoarchitectonic atlas-based interrogation of fMRI data sets. Neuroimage 19 (3), 1233–1239.

Martin, L.N., Delgado, M.R., 2011. The influence of emotion regulation on decision-making under risk. J. Cogn. Neurosci. 23 (9), 2569–2581.

Matsumoto, K., Suzuki, W., Tanaka, K., 2003. Neuronal correlates of goal-based motor selection in the prefrontal cortex. Science 301 (5630), 229–232.

Mazaika, P., Hoeft, F., Glover, G.H., Reiss, A.L., 2009. Methods and software for fMRI analysis for clinical subjects. Poster at Hum. Brain Mapp.

McClure, S.M., Laibson, D.I., Loewenstein, G., Cohen, J.D., 2004. Separate neural systems value immediate and delayed monetary rewards. Science 306 (5695), 503–507.

McDannald, M. a, Lucantonio, F., Burke, K. a, Niv, Y., Schoenbaum, G., 2011. Ventral striatum and orbitofrontal cortex are both required for model-based, but not model-free, reinforcement learning. J. Neurosci. 31 (7), 2700–2705.

Mechsner, F., Kerzel, D., Knoblich, G., Prinz, W., 2001. Perceptual basis of bimanual coordination. Nature 414 (6859), 69–73.

Nichols, T., Brett, M., Andersson, J., Wager, T., Poline, J.-B., 2005. Valid conjunction inference with the minimum statistic. Neuroimage 25 (3), 653–660.

Nieuwenhuis, S., Forstmann, B.U., Wagenmakers, E.-J., 2011. Erroneous analyses of interactions in neuroscience: a problem of significance. Nat. Neurosci. 14 (9), 1105–1107.

Noonan, M.P., Walton, M.E., Behrens, T.E.J., Sallet, J., Buckley, M.J., Rushworth, M.F.S., 2010. Separate value comparison and learning mechanisms in macaque medial and lateral orbitofrontal cortex. Proc. Natl. Acad. Sci. U.S.A. 107 (47), 20547–20552.

O'Doherty, J.P., Deichmann, R., Critchley, H.D., Dolan, R.J., 2002. Neural responses during anticipation of a primary taste reward. Neuron 33 (5), 815–826.

O'Reilly, R.C., Frank, M.J., 2006. Making working memory work: a computational model of learning in the prefrontal cortex and basal ganglia. Neural Comput. 18 (2), 283–328.

O'Reilly, J.X., Woolrich, M.W., Behrens, T.E.J., Smith, S.M., Johansen-Berg, H., 2012. Tools of the trade: psychophysiological interactions and functional connectivity. Soc. Cogn. Affect. Neurosci. 7 (5), 604–609 (Jun).

Oliveira, F.T.P., Ivry, R.B., 2008. The representation of action: insights from bimanual coordination. Curr. Dir. Psychol. Sci. 17 (2), 130–135.

Padoa-Schioppa, C., Assad, J.a., 2006. Neurons in the orbitofrontal cortex encode economic value. Nature 441 (7090), 223–226.

Palminteri, S., Boraud, T., Lafargue, G., Dubois, B., Pessiglione, M., 2009. Brain hemispheres selectively track the expected value of contralateral options. J. Neurosci. 29 (43), 13465–13472.

Penfield, W., Jasper, H., 1954. Epilepsy and the Functional Anatomy of the Human Brain. Little Brown, Boston.

Penny, W.D., Henson, R., 2006. Analysis of Variance. In: Friston, K.J., Ashburner, J.T., Kiebel, S.J., Nichols, T.E., Penny, W.D. (Eds.), Statistical Parametric Mapping: The Analysis of Functional Brain Images. Academic Press, London, pp. 166–177 (Ch. 13).

Pesaran, B., Nelson, M.J., Andersen, R.A., 2006. Dorsal premotor neurons encode the relative position of the hand, eye, and goal during reach planning. Neuron 51 (1), 125–134.

Peters, J., Büchel, C., 2009. Overlapping and distinct neural systems code for subjective value during intertemporal and risky decision making. J. Neurosci. 29 (50), 15727–15734.

Phelps, E.A., Delgado, M.R., Nearing, K.I., LeDoux, J.E., 2004. Extinction learning in humans: role of the amygdala and vmPFC. Neuron 43 (6), 897–905.

Picard, N., Strick, P.L., 2001. Imaging the premotor areas. Curr. Opin. Neurobiol. 11 (6), 663–672.

Plassmann, H., O'Doherty, J., Rangel, A., 2007. Orbitofrontal cortex encodes willingness to pay in everyday economic transactions. J. Neurosci. 27 (37), 9984–9988.

Plassmann, H., O'Doherty, J., Shiv, B., Rangel, A., 2008. Marketing actions can modulate neural representations of experienced pleasantness. Proc. Natl. Acad. Sci. U.S.A. 105 (3), 1050–1054 (Jan).

Platt, M.L., Glimcher, P.W., 1999. Neural correlates of decision variables in parietal cortex. Nature 400 (6741), 233–238.

Platt, M., Padoa-Schioppa, C., 2008. Neuronal representations of value. In: Glimcher, P.W., Camerer, C.F., Fehr, E., Poldrack, R.A. (Eds.), Neuroeconomics: Decision Making and the Brain. Academic Press, San Diego, CA, pp. 441–462 (Ch. 29).

Poldrack, R.A., Clark, J., Paré-Blagoev, E.J., Shohamy, D., Creso Moyano, J., Myers, C., Gluck, M.A., 2001. Interactive memory systems in the human brain. Nature 414 (6863), 546–550.

Poldrack, R.A., Fletcher, P.C., Henson, R.N., Worsley, K.J., Brett, M., Nichols, T.E., 2008. Guidelines for reporting an fMRI study. Neuroimage 40 (2), 409–414.

Rescorla, R.A., Wagner, A.R., 1972. A theory of Pavlovian conditioning: variations in the effectiveness of reinforcement and non-reinforcement. In: Black, A.H., Prokasy, W.F. (Eds.), Classical Conditioning II: Current Research and Theory. Appleton-Century-Crofts, pp. 64–99.

Rosenbaum, D.A., Dawson, A.M., Challis, J.H., 2006. Haptic tracking permits bimanual independence. J. Exp. Psychol. Hum. Percept. Perform. 32 (5), 1266–1275.

Rudebeck, P.H., Murray, E.A., 2011. Dissociable effects of subtotal lesions within the macaque orbital prefrontal cortex on reward-guided behavior. J. Neurosci. 31 (29), 10569–10578.

Rudebeck, P.H., Behrens, T.E., Kennerley, S.W., Baxter, M.G., Buckley, M.J., Walton, M.E., Rushworth, M.F.S., 2008. Frontal cortex subregions play distinct roles in choices between actions and stimuli. J. Neurosci. 28 (51), 13775–13785.

Sadato, N., Yonekura, Y., Waki, A., Yamada, H., Ishii, Y., 1997. Role of the supplementary motor area and the right premotor cortex in the coordination of bimanual finger movements. J. Neurosci. 17 (24), 9667–9674.

Scherberger, H., Andersen, R.a., 2007. Target selection signals for arm reaching in the posterior parietal cortex. J. Neurosci. 27 (8), 2001–2012.

Schlag, J., Schlag-Rey, M., 1987. Evidence for a supplementary eye field. J. Neurophysiol. 57 (1), 179–200.

Schönberg, T., Daw, N.D., Joel, D., O'Doherty, J.P., 2007. Reinforcement learning signals in the human striatum distinguish learners from nonlearners during reward-based decision making. J. Neurosci. 27 (47), 12860–12867.

Schwarz, G., 1978. Estimating the dimension of a model. Ann. Stat. 6 (2), 461–464.

Seo, H., Lee, D., 2007. Temporal filtering of reward signals in the dorsal anterior cingulate cortex during a mixed-strategy game. J. Neurosci. 27 (31), 8366–8377.

Shackman, A.J., Salomons, T.V., Slagter, H.a., Fox, A.S., Winter, J.J., Davidson, R.J., 2011. The integration of negative affect, pain and cognitive control in the cingulate cortex. Nat. Rev. Neurosci. 12 (3), 154–167.

Shadlen, M.N., Newsome, W.T., 2001. Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. J. Neurophysiol. 86 (4), 1916–1936.

Shadmehr, R., 1997. Neural correlates of motor memory consolidation. Science 277 (5327), 821–825.

Simon, D.A., Daw, N.D., 2011a. Environmental statistics and the trade-off between model-based and TD learning in humans. In: Shawe-Taylor, J., Zemel, R.S., Bartlett, P., Pereira, F.C.N., Weinberger, K.Q. (Eds.), Advances in Neural Information Processing Systems, 24, pp. 127–135.

Simon, D.A., Daw, N.D., 2011b. Neural correlates of forward planning in a spatial decision task in humans. J. Neurosci. 31 (14), 5526–5539.

Snyder, L., Batista, A., Andersen, R., 1997. Coding of intention in the posterior parietal cortex. Nature 386 (6621), 167–170.

Sperry, R.W., 1961. Cerebral organization and behavior: the split brain behaves in many respects like two separate brains, providing new research possibilities. Science 133 (3466), 1749–1757.

Stephan, K.E., Penny, W.D., Daunizeau, J., Moran, R.J., Friston, K.J., 2009. Bayesian model selection for group studies. Neuroimage 46 (4), 1004–1017.

Sugrue, L., Corrado, G., 2004. Matching behavior and the representation of value in the parietal cortex. Science 304 (June), 1782–1787.

Swinnen, S., Wenderoth, N., 2004. Two hands, one brain: cognitive neuroscience of bimanual skill. Trends Cogn. Sci. 8 (1), 18–25.

Tanji, J., Okano, K., Sato, K.C., 1988. Neuronal activity in cortical motor areas related to ipsilateral, contralateral, and bilateral digit movements of the monkey. J. Neurophysiol. 60 (1), 325–343.

Tom, S.M., Fox, C.R., Trepel, C., Poldrack, R.A., 2007. The neural basis of loss aversion in decision-making under risk. Science 315 (5811), 515–518.

Tzourio-Mazoyer, N., Landeau, B., Papathanassiou, D., Crivello, F., Etard, O., Delcroix, N., Mazoyer, B., Joliot, M., 2002. Automated anatomical labeling of activations in SPM using a macroscopic anatomical parcellation of the MNI MRI single-subject brain. Neuroimage 15 (1), 273–289.

Venkatraman, V., Rosati, A.G., Taren, A.A., Huettel, S.A., 2009. Resolving response, decision, and strategic control: evidence for a functional topography in dorsomedial prefrontal cortex. J. Neurosci. 29 (42), 13158–13164.

Wallis, J.D., 2012. Cross-species studies of orbitofrontal cortex and value-based decision-making. Nat. Neurosci. 15 (1), 13–19.

Wallis, J.D., Miller, E.K., 2003. Neuronal activity in primate dorsolateral and orbital prefrontal cortex during performance of a reward preference task. Eur. J. Neurosci. 18 (7), 2069–2081.

Walton, M.E., Behrens, T.E.J., Buckley, M.J., Rudebeck, P.H., Rushworth, M.F.S., 2010. Separable learning systems in the macaque brain and the role of orbitofrontal cortex in contingent learning. Neuron 65 (6), 927–939.

Wimmer, G.E., Daw, N.D., Shohamy, D., 2012. Generalization of value in reinforcement learning by humans. Eur. J. Neurosci. 35 (7), 1092–1104.

Wunderlich, K., Rangel, A., O'Doherty, J.P., 2009. Neural computations underlying action-based decision making in the human brain. Proc. Natl. Acad. Sci. U.S.A. 106 (40), 17199–17204.

Wunderlich, K., Rangel, A., O'Doherty, J.P., 2010. Economic choices can be made using only stimulus values. Proc. Natl. Acad. Sci. U.S.A. 107 (34), 15005–15010.

Xue, G., Ghahremani, D.G., Poldrack, R.A., 2008. Neural substrates for reversing stimulus–outcome and stimulus–response associations. J. Neurosci. 28 (44), 11196–11204.

Yu, C., Zhou, Y., Liu, Y., Jiang, T., Dong, H., Zhang, Y., Walter, M., 2011. Functional segregation of the human cingulate cortex is confirmed by functional connectivity based neuroanatomical parcellation. Neuroimage 54 (4), 2571–2581.