Title
Differentiating neural systems mediating the acquisition vs. expression of goal-directed and habitual behavioral control

Journal
European Journal of Neuroscience, 41(10)

ISSN
0953-816X

Authors
Liljeholm, Mimi
Dunne, Simon
O'Doherty, John P

Publication Date
2015-05-01

DOI
10.1111/ejn.12897

Peer reviewed

HHS Public Access

Author manuscript

Eur J Neurosci. Author manuscript; available in PMC 2016 May 01.

Published in final edited form as:

Eur J Neurosci. 2015 May ; 41(10): 1358–1371. doi:10.1111/ejn.12897.

Differentiating neural systems mediating the acquisition versus expression of goal-directed and habitual behavioral control

Mimi Liljeholm[1,2], **Simon Dunne**[1,3], and **John P. O'Doherty**[1,3]

[1]Trinity College Institute of Neuroscience, Trinity College Dublin, Dublin 2, Ireland [2]Department of Cognitive Sciences, University of California, Irvine, 92697 [3]Division of the Humanities and Social Sciences and Computation and Neural Systems Program, California Institute of Technology, Pasadena, California 91125

Abstract

Considerable behavioral data indicates that operant actions can become habitual, as evidenced by insensitivity to changes in the action-outcome contingency and in subjective outcome values. Notably, although several studies have investigated the neural substrates of habits, none has clearly differentiated the areas of the human brain that support habit formation from those that implement habitual control. We scanned participants with fMRI as they learned and performed an operant task in which the conditional structure of the environment encouraged either goal-directed encoding of the consequences of actions, or a habit-like mapping of actions to antecedent cues. Participants were also scanned during a subsequent assessment of insensitivity to outcome devaluation. We identified dissociable roles of the cerebellum and ventral striatum, across learning and test performance, in behavioral insensitivity to outcome devaluation. We also show that the inferior parietal lobule – an area previously implicated in several aspects of goal-directed action selection, including the attribution of intent and awareness of agency – predicts sensitivity to outcome devaluation. Finally, we reveal a potential functional homology between the human subgenual cortex and rodent infralimbic cortex in the implementation of habitual control. In summary, our findings suggest a broad systems division, at the cortical and subcortical levels, between brain areas mediating the encoding and expression of action-outcome and stimulus-response associations.

Corresponding author: Mimi Liljeholm, Department of Cognitive Sciences, 2312 Social & Behavioral Science Gateway Building, University of California, Irvine, California 92697, Phone: 949-824-8444, Fax: 949-824-2307, m.liljeholm@uci.edu.

There are no conflicts of interest.

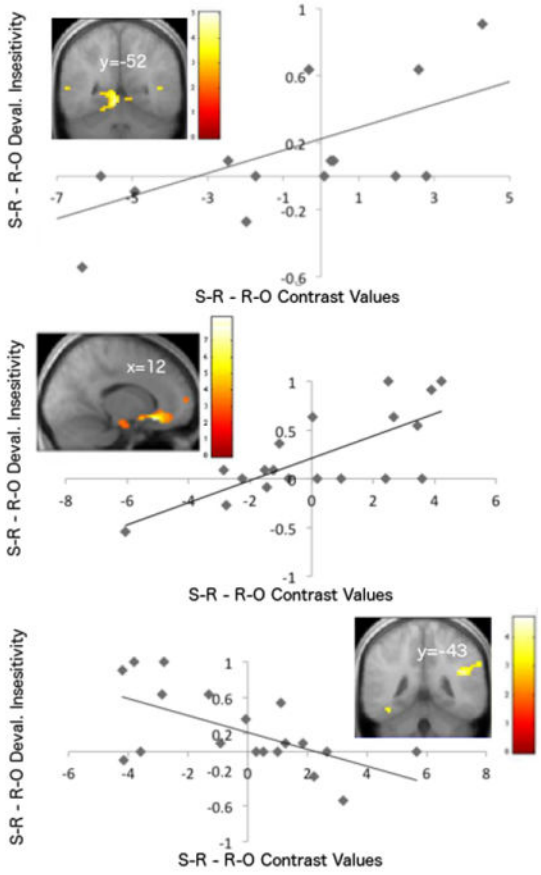

Keywords

Habit formation; stimulus-response; action-outcome; fMRI

Introduction

Habitual action selection is defined by insensitivity to changes in the causal efficacy with which actions produce rewards and to the current subjective value of those rewards (Balleine and Dickinson, 1998). Neuroscientific research on humans and rodents has demonstrated that brain areas mediating habitual performance are dissociable from those supporting more deliberate, goal-directed, action selection (Balleine and Dickinson, 1998; Yin et al., 2004, 2005; Valentin et al., 2007; Tricomi et al., 2009). Intriguingly, work in rodents also suggests that distinct neural substrates make specialized contributions to the development versus deployment of habits (Killcross and Coutureau, 2003). In contrast, in humans there has been no clear differentiation between brain areas that support habit formation and those that implement habitual control. Although a couple of neuroimaging studies have demonstrated increases in neural activity concomitant with the development of habits in a posterior area of the lateral striatum (Tricomi et al., 2009; Wunderlich et al., 2012), the use of overtraining to induce habitual responding in these studies confounds well-trained performance with

habitual control. In the current study, in order to discriminate between neural substrates supporting the acquisition versus expression of habits, we scanned human participants with functional magnetic resonance imaging (fMRI) as they performed a novel instrumental task (see Task in Methods), designed to rapidly induce habitual responding without the potentially confounding process of overtraining.

Pharmacological disruptions and electrophysiological recordings of the rodent brain have strongly implicated the dorsolateral striatum (DLS) in the acquisition of habits: Pre-training lesions of the DLS abolish habit formation (Yin et al., 2004), and distinct changes in neuronal activity patterns (Jog et al., 1999), including substantial decreases in firing rates (Carelli et al., 1997; Tang et al., 2007), have been demonstrated in this area across the development of habitual responding. In contrast, the infralimbic region of the prefrontal cortex has been suggested to support an executive control system that facilitates the expression of habits: Post-training muscimol induced inactivation of the infralimbic cortex disrupts habitual performance (Coutureau and Killcross, 2003; Haddon and Killcross, 2011), and changes in neuronal ensemble activity patterns in this area occur very late in training and closely track the behavioral manifestation of habitual control (Smith and Graybiel, 2013). Further evidence for the involvement of the infralimbic cortex in the expression of habits comes from studies in which optogenetic perturbation of this area disrupts well-ingrained habitual behavior (Smith et al., 2012). Based on these findings in rodents, we hypothesized that human homologues of the DLS and infralimbic cortex would be involved in the formation and expression of habits respectively, such that behavioral insensitivity to outcome devaluation would correlate with neural activity during learning in the dorsal putamen (Carelli and West, 1991; Draganski et al., 2008), but with activity during actual test performance in the subgenual cortex.(Ongur and Price, 2000; Ongur et al., 2003).

Methods

Participants

Twenty volunteers (mean age =21.4 ± 2.63, range=19-28; 11 males) participated in the experiment. Due to technical problems (loss of power to the stimulus computer), one of the subjects was excluded, yielding a total of 19 participants. All participants were normal and healthy, and were recruited locally from the city of Dublin, Ireland. The study was approved by the Trinity College Dublin School of Psychology research ethics committee, and all participants gave informed consent. The study conformed to the guidelines set out in the 2013 WMA Declaration of Helsinki.

Task

Goal-directed actions, defined by their sensitivity to changes in both action-outcome contingency and outcome value, have been proposed to depend on an internal model of the world that explicitly relates alternative actions to future environmental states (Doya et al., 2002; Daw et al., 2005). Consistent with this theoretical framework, data from rodent studies suggests that reliance on a goal-directed versus habitual strategy might depend on the ease with which alternative actions can be associated with distinct outcomes: goal-directed performance appears to dominate, in spite of overtraining, when alternative actions yield

distinct sensory-specific outcomes (e.g., grain versus sucrose pellets) (Colwill and Rescorla, 1985; Holland, 2004), as well as when the rate of outcome delivery depends on the rate of responding, rather than on a particular time interval passing between successive reinforced responses (i.e., ratio versus interval schedules of reinforcement) (Dickinson, 1983). The current task is structured on these potential bases of behavioral control, in that external contingencies either facilitate or impede a reliable mapping of alternative actions to distinct sensory-specific outcome states.

Participants were required to maintain the balance of a system of fluid-filled beakers (see Fig. 1A for details). As long as all beakers had sufficient fluid, system balance was maintained and randomly occurring balance checks yielded monetary reward. However, on each trial, one of the beakers would be emptied causing “system imbalance”, with balance checks resulting in monetary loss until the participant re-filled the beaker by performing a particular instrumental action. The emptying of a beaker was always accompanied by the onset of one of four abstract cues. In the *Stimulus-Response (S-R)* condition, the identity of the presented cue determined which instrumental action would re-fill the emptied beaker, regardless of *which* of the beakers had lost its fluid. Consequently, across trials, each rewarded action was paired with a specific antecedent cue, but was decorrelated from the re-filling of any particular beaker. Conversely, in a *Response-Outcome (R-O)* condition, each instrumental action re-filled a particular beaker, regardless of which abstract cue was presented, such that identification of the relevant subgoal (e.g., re-filling beaker 1), combined with knowledge about specific action-outcome contingencies (i.e., action 1 re-fills beaker 1), indicated which action would restore system balance. To ensure that discriminatory neural activity was not due to differences in the visual processing of abstract cues versus beakers, a matching task was interleaved with the instrumental task (see Fig. 1B).

The persistent execution of an action after its outcome has been devalued is a defining feature of habitual performance (Adams and Dickinson, 1981; Adams, 1982). In the current study, following acquisition of the instrumental task, we devalued one of the four beakers by degrading its relationship to monetary gain, such that system balance was maintained, and continued to yield points, even when the liquid in this beaker dropped below threshold. Because there was a small cost for regulating the system, attempts to re-fill this beaker now resulted in a net loss (see Fig. 1A and Experimental Procedure for details). Based on the notion that failure to associate alternative actions with distinct outcome states obstructs goal-directed encoding, we predicted that the S-R condition would bolster habit formation during acquisition and bias participants towards habitual action selection, defined as responding to re-fill the devalued beaker, in subsequent tests performance. Note that, in both conditions, distinct features of the stimulus environment can enter into stimulus-response associations; whereas in the S-R condition, each rewarded action is reliably preceded by a particular arbitrary cue, rewarded actions in the R-O condition are reliably preceded by the emptying of a particular beaker. However, critically, it is only in the R-O condition that alternative actions can be associated with distinct outcome states. Consequently, we expected performance in this condition to be goal-directed.

Experimental Procedure

We scanned participants as they acquired and performed the instrumental task, as well as during a subsequent devaluation test phase. Each subject participated in both the R-O and S-R condition; with conditions being ran in separate, immediately consecutive, sessions (order counterbalanced across subjects) and with a novel set of four instrumental actions being used in each condition. Each condition included a response pre-training phase, a learning phase, a devaluation phase, and a final test phase as described below. The response-training phase of the first condition was conducted outside the scanner (in a separate testing room), before participants were transferred to the scanner in which they remained throughout all subsequent stages of the experiment. Before being transferred to the scanner, participants were presented with a cover-story describing the beaker system and the task. They were told that they would be in one of two possible conditions – one in which each instrumental action re-filled a particular beaker regardless of which cue was presented, and another in which the identity of the cue determined which of the four actions was required to re-fill an emptied beaker, regardless of the identity of that beaker – and that part of their task was to determine which of the two conditions they were in. The entire experiment lasted for approximately 2 hours, with 1.5 hours being spent in the scanner, and with approximately 60 minutes of active scanning, during the learning phases and devaluation tests in each condition.

Response pre-training—Prior to the instrumental learning phases, participants received pre-training on the four instrumental actions (each a 3-press sequence). During this training, key-press sequences were illustrated by a white dot moving across three gray squares, horizontally aligned at the center of the screen. Initially, participants viewed and then immediately attempted to replicate each sequence, with feedback (i.e., correct/incorrect) given on each trial. After a total of 5 correct replications of each response-sequence, they proceeded to a retrieval phase, in which they had to generate each unique sequence at least 5 times without any prompts, again with feedback given at the completion of each sequence. Participants were allowed to repeat these two phases as many times as they wanted to, knowing that they would have to use the actions to earn monetary reward in a subsequent phase.

Instrumental learning phase—The instrumental task was as illustrated in Figure 1A. Note that, in addition to the increase or decrease in monetary points based on system balance, there was a small cost for regulating the system. This response cost was included to ensure that, during test, participants would not respond simply based on any reinforcement intrinsic to executing the correct response. The response cost message (screen 3 in Fig. 1A) also served to inform participants that their action had successfully regulated the system, rather than the system having self-regulated (in which case no response cost was charged) due to failure to perform the correct action within 7 seconds of trial onset. Participants were allowed to perform as many key-presses as they wanted during system imbalance (i.e., during the 7 seconds following trial onset), with the correct sequence of three consecutive key-presses immediately restoring system balance and terminating the trial. There was no constraint on the temporal spacing between key-presses, as long all three presses were performed consecutively within the trial window.

Critically, the stimulus materials presented in Figure 1A were identical across the two conditions: our manipulation consisted entirely of differences in the contingencies between cues, actions and beaker outcomes. To rule out visual processes involved in selectively attending to the abstract cues versus the beakers as a source of any imaging effects, a matching task was block-interleaved with the instrumental task during instrumental learning (see Fig. 1B). Briefly, in matching blocks, the inter-trial intervals and trial onsets were exactly as in the system balance task, except that the words “Matching trial” were displayed center screen: Without this indication, matching trials would be identical to instrumental trials during the relevant trial period, and thus could not have served as controls. Following the appearance of the abstract cue and emptying of the relevant beaker, a white masking screen was displayed, followed by a depiction of either an abstract cue (S-R condition) or a set of beakers (R-O condition), together with a query about whether the currently shown cue/beaker set matched that on the previous screen. In each condition, the instrumental learning phase consisted of 4 blocks of trials, with each block being further divided into one sub-block of 24 instrumental trials, followed by a sub-block of 8 matching trials, and with the order of trials randomized within each sub-block.

At the end of the instrumental learning phase, participants were asked whether they felt confident that they had learned how to regulate the system or whether they wanted to receive an additional set of 20 training trials (5 with each action). Five participants (two in the R-O and three in the S-R condition) requested and received additional training. No scanning was conducted during additional training, nor were the added trials included in assessments of accuracy and response times during acquisition.

Devaluation phase—Following instrumental learning, participants were instructed that the system had changed such that one of the beakers was no longer relevant for system balance, which would be maintained, and continue to yield points, even when the liquid in this beaker dropped below threshold. They then observed as the system regulated itself (i.e., no actions were performed) across 16 trials (4 with each beaker) in order to discover the identity of the devalued beaker. Participants were told that they would not lose or gain any of the displayed points during this phase. In the imaging experiment, two participants failed to correctly identify the devalued beaker after this devaluation procedure, which was therefore repeated once for these participants. All other participants successfully identified the devalued beaker at the end of the devaluation procedure.

Test phase—Having correctly identified the devalued beaker, participants were again given the opportunity to regulate the system for personal monetary reward. During this phase, all text messages indicating gains or losses were covered up, in order to prevent additional learning (i.e., simulating extinction). Participants were instructed that, in spite of these gray strips, they should assume that all was exactly as they had learned before: that is, they would still lose points whenever the system was not balanced, there was still a cost for regulating the system, and the previously identified irrelevant beaker was still irrelevant for system balance. Importantly, because there was a small charge for each instrumental regulation of the system, re-filling the now irrelevant beaker resulted in a net monetary loss. The test phase consisted of a single instrumental block with 44 randomly ordered trials; 11

trials with each beaker, including the devalued **(i.e., irrelevant)** one. We operationally define devaluation insensitivity, our assay of habitual performance, as the proportion of the 11 devalued trials on which participants initiated a response to refill the beaker. Thus, if a participant initiated a response on 2 of those 11 trials, their devaluation insensitivity score would be 0.18.

Pay-off structure—Throughout the instrumental task, system balance checks occurred on average every 3 seconds. During the ITI (mean=5sec.), the system was always balanced, with each check yielding a reward of 0.5 points. At each trial onset, the system became imbalanced, with each balance check resulting in a loss of 0.5 points, and would remain so for 7 seconds unless regulated by the participant. Thus, if no action were taken to balance the system, each trial would entail an average loss of 1.17 points. If the system was balanced immediately, no points were lost due to system imbalance, but there was a response cost of 0.1 points (this cost was only applied to correct responses resulting in system regulation). Thus, the average gain of responding to balance the system was 1.07 points. After beaker devaluation, during the test phase, this was still the case for regulatory responses re-filling non-devalued beakers, however, responding to re-fill the devalued beaker now resulted in a net loss of 0.1 points (the response cost).

Imaging acquisition and analyses

A 3 Tesla scanner (MAGNETOM Trio; Siemens) was used to acquire structural T1-weighted images and T2*-weighted echoplanar images (repetition time = 2.65 s; echo time = 30 ms; flip angle = 90°; 45 transverse slices; matrix = 64 × 64; field of view = 192 mm; thickness = 3 mm; slice gap = 0 mm) with BOLD contrast. To recover signal loss from dropout in the medial orbitofrontal cortex (mOFC)(O'Doherty et al., 2002), each horizontal section was acquired at 30° to the anterior commissure–posterior commissure axis. Image processing and statistical analyses were performed using SPM8 (http://www.fil.ion.ucl.ac.uk/spm). The first four volumes of images were discarded to avoid T1 equilibrium effects. All remaining volumes were corrected for differences in slice acquisition, realigned to the first volume, spatially normalized to the Montreal Neurological Institute (MNI) echoplanar imaging template, and spatially smoothed with a Gaussian kernel (8 mm, full width at half-maximum). We used high-pass filter with a cutoff of 128 s.

As previously noted (see Experimental Procedure), instrumental conditions were ran in separate consecutive sessions, with the acquisition phase of each condition consisting of interleaved instrumental and matching block, and with active scanning occurring only during the acquisition and test phase of each condition (i.e., the scanner was off during response pre-training and devaluation phases), for a total of 4 active scanning periods per subject. For each subject, we constructed an fMRI design matrix that included the acquisition and test phases of both experimental conditions. For each of the two instrumental acquisition phases (i.e., R-O and S-R), stick functions modeled the trial onsets in each block of instrumental learning and in matching blocks. In addition, three regressors, respectively modeling the onsets of error trials, the times of all individual key-presses and the times of all point displays (i.e., gains and losses) were entered as regressors of no interest, together with six regressors accounting for the residual effects of head motion. For each of the two test

phases, two stick functions respectively modeled non-devalued and devalued trials: again, for each instrumental condition, additional regressors modeling the onsets of error trials and the times of all key-presses were added together with six regressors accounting for the residual effects of head motion as regressors of no interest. All regressors were convolved with a canonical hemo-dynamic response function.

Group-level random-effects statistics were generated by entering contrasts of parameter estimates for the different regressors into between-subjects analyses of variance (ANOVA). Specifically, to delineate neural substrates engaged during early acquisition versus during expression, contrasts of parameter estimates for the first two blocks of instrumental learning, and for the devaluation test phase, were entered for each instrumental condition into a 2×2 (condition by experimental phase) ANOVA. Contrasts of parameter estimates for all learning blocks and for matching blocks were entered into a separate 2× 2(stimulus by task) ANOVA, to compare differences between instrumental conditions during learning to those between matching control conditions. An analogous ANOVA was performed using contrasts of parameter estimates for devaluation test trials and matching trials. Finally, a 4×2 (instrumental by acquisition block) ANOVA assessed training-related changes in neural activity, with contrasts of parameter estimates for each of the four blocks of instrumental learning, for each instrumental condition. Following estimation of the 2^{nd} level model, F-tests were specified by adding linear weights to each block of learning (e.g., [-1.5 -0.5 0.5 1.5] for increases across blocks and [1.5 0.5 -0.5 -1.5] for decreases), in each instrumental condition.

To assess whether neural discrimination between instrumental conditions correlated with the degree of devaluation insensitivity, a simple t-tests was performed on first level interaction contrasts [SR>RO × Instrumental>Matching] of parameter estimates from the learning phase, with the degree of devaluation insensitivity entered as a covariate. An analogous t-test using interaction contrasts of parameter estimates from the devaluation test phase was also performed, and exclusive functional masks were used to assess the specificity of neural effects, such that all voxels that reached significance at a threshold of $p<0.1$ when assessing correlates of devaluation insensitivity during the test phase were removed from the effects observed during learning, and vice versa. Exclusive functional masks were also used to assess the directions of simple effects underlying the Instrumental Condition by Experimental Phase interaction. Finally, we used exclusive functional masks to selectively assess training-related increases versus decreases in neural activity, such that, for example, when testing for increases in neural activity across blocks of training in the R-O condition, voxels were removed that reached significance at $p<0.1$ for the same test in the S-R condition.

Small volume corrections (SVC) were performed on several a priori regions of interest using a 10 mm sphere, with center coordinates obtained from highly relevant studies. Specifically, in an analogous study on observational learning (Liljeholm et al., 2012), we found selective recruitment of the extra-striate cortex (±45, -72, 9), the tail of caudate nucleus/thalamus (±21, -30, 9) and lingual gyrus (±12, -69, 0) by the S-R condition during acquisition. Conversely, areas selective recruited by the R-O condition in that same observational learning study, and also identified in our other work on goal-directed performance

(Liljeholm et al., 2011), include the inferior parietal lobule (±51, -52, 33) and anterior caudate nucleus (±16, 8, 19). Finally, several human neuroimaging studies have implicated a posterior region of the putamen (±33, -24, 0) in habit formation (Tricomi et al., 2009; de Wit et al., 2012; Wunderlich et al., 2012; Lee et al., 2014), and the ventromedial prefrontal cortex (±4, 55, -7) in goal-directed processes (Hampton et al., 2006; Tanaka et al., 2008; Liljeholm et al., 2011). It should be noted that, although specified a priori based on a closely related literature, effects in several of these regions survived whole-brain cluster-size threshold (CST) correction, as well as SVC, in the current study.

As noted, the putamen, considered a human homolog of the rodent DLS based on its afferent and efferent projections (Carelli and West, 1991; Draganski et al., 2008), has been implicated in habitual performance in several human imaging studies (Tricomi et al., 2009; de Wit et al., 2012; Wunderlich et al., 2012; Lee et al., 2014). In contrast, the role of infralimbic cortex in instrumental performance has been exclusively demonstrated in rodents (Killcross and Coutureau, 2003; Haddon and Killcross, 2011; Smith et al., 2012; Smith and Graybiel, 2013). Consequently, and based on anatomical and functional evidence for a homology between the rodent infralimbic cortex and the human subgenual cortex (Ongur and Price, 2000; Ongur et al., 2003; Drevets et al., 2008), predictions regarding this area where tested using an anatomical masks of the subgenual cortex (BA25), defined using the Wake Forest University (WFU) Pickatlas (Maldjian et al., 2003). All other effects were reported at $p < 0.05$, using cluster size thresholding (CST) of SPM t-maps to adjust for multiple comparisons (Forman et al., 1995). AlphaSim, a Monte Carlo simulation (AFNI) was used to determine cluster size and significance. Using an individual voxel probability threshold of $p=0.005$ indicated that using a minimum cluster size of 111 MNI transformed voxels resulted in an overall significance of $p < 0.05$.

To eliminate non-independence bias for plots of contrast values, a leave-one-subject-out (LOSO (Esterman et al., 2010) approach was used, in which 19 GLMs were run with one subject left out in each, and with each GLM defining the voxel cluster for the left out subject. Using rfxplot (Glascher, 2009), mean contrast values were extracted from spheres (10 mm) centered on these LOSO peaks (identified within ROIs for small volume corrections) and were averaged across subjects to plot overall effect sizes.

Results

Behavioral results

The results from the test phase indicated that our manipulation did indeed produce differences in devaluation sensitivity: Having correctly identified the devalued beaker, participants initiated a response on a significantly greater proportion of devalued trials in the S-R condition (mean proportion = 0.43, SEM=0.09) than in the R-O (mean proportion = 0.19, SEM=0.09), $t(18)=2.3$, $p<0.031$. Indeed, whereas only 5 of 19 participants initiated a response on any devalued trials in the R-O condition, 15 of 19 participants initiated a response on at least one devalued trial in the S-R condition; $\chi^2=10.56$, $p<0.005$.

Note that devaluation insensitivity was defined such that a response to obtain the devalued outcome (i.e., to re-fill the devalued beaker) was counted even if the entire three-press

sequence was not completed on that trial. Indeed, in the S-R condition, when a participant responded to fill up the devalued beaker, they often did so without completing the entire three-press sequence, consistent with evidence from the rodent literature that the response most proximal to reward remains sensitive to devaluation long after more distal responses have become habitual (Killcross & Coutureau, 2003). In contrast, those few participants that responded on devalued trials in the R-O condition did so on most or all devalued trials and completed all three key-presses whenever initiating a response, suggesting a more deliberate decision to respond.

Importantly, during the test phase, whereas in the R-O condition participants could determine both the accuracy and value of a particular action solely by identifying the emptied beaker, in the S-R condition, determining the value and accuracy of a given response required identification of both the abstract cue and the emptied beaker. It is possible therefore, that differences in devaluation performance were due to this additional aspect of the S-R condition. The overall pattern of behavioral results, however, strongly suggests that this was not the case, and generally rules out task difficulty as the source of our effects. First, there were no significant differences between conditions in the percent of incorrect responses, during either acquisition (R-O mean=14%, SEM=3, S-R mean=11%, SEM=2; t(18)=1.1, p=0.30) or test performance (R-O mean=4%, SEM=1, S-R mean=4%, SEM=2; t(18)=0, p=1.0). Likewise, response times did not differ between conditions during either acquisition (R-O mean=1482 ms, SEM=62; S-R mean=1393 ms, SEM=78; t(18)=1.95, p=0.21) or test performance (R-O mean=1258 ms, SEM=54, S-R mean=1361 ms, SEM=68; t(18)=-1.52, p=0.15). The fact that there were no significant differences between conditions in either accuracy or reaction times during test makes it highly unlikely that differences in devaluation sensitivity reflected differences in task difficulty. Perhaps most pertinently, differences in devaluation insensitivity were not correlated with differences in accuracy (p=0.5) or reaction times (p=0.8), nor did individual differences in accuracy or reaction times predict neural effects in any of the areas identified by our imaging analyses.

To assess the influence of counterbalancing order, we performed order-by-condition analyses of variance on response times and accuracy scores during the learning and test phases, as well as on the devaluation insensitivity scores. There was no significant interaction between order and condition for devaluation insensitivity ($F(1,17)=0.55$, p=0.47), nor for either response times ($F(1,17)=0.89$, p=0.36) or accuracy scores ($F(1,17)=0.07$, p=0.79) during the test phase. In contrast, during training, there was an anticipated interaction for both response times ($F(1,17)=5.96$, $p<0.03$) and accuracy scores ($F(1,17)=23.01$, $p<0.001$), such that response times were longer and the percent incorrect was greater for whatever condition came first: Thus, while the difference in response times between instrumental conditions differed across the two orders (mean difference for R-O first=0.19; mean difference for S-R first=-0.32) there was no difference between instrumental conditions when compared for a given order (i.e., for R-O vs. S-R with both first, p=0.84 and for R-O vs. S-R with both second, p=0.71). A similar pattern was observed for percent incorrect scores (mean difference for R-O first=0.08; mean difference for S-R first=-0.01; significance test for R-O vs. S-R with both first, p=0.26 and for R-O vs. S-R

with both second, p=0.90). Consequently, we can rule out counterbalancing order as a source of our behavioral and imaging effects of interest.

There were no differences between instrumental conditions in the number of requested repetitions of response pre-training phases (R-O mean=1.98, SEM=0.29; S-R mean=1.75, SEM=0.18; t(18)=0.72, p=0.48), nor in the total number of key presses executed during instrumental learning (R-O mean=470.63, SEM=32.14; S-R mean=503.74, SEM=31.63; t(18)=-0.85, p=0.40) or on non-devalued trials during test (R-O mean=121.84, SEM=4.57; S-R mean=125.53, SEM=4.53; t(18)=-1.34, p=0.57). Finally, there was no difference between instrumental conditions in the total number of points earned during instrumental learning (R-O mean=33.6, SEM=3.61; S-R mean=33.4, SEM=3.33; t(18)=0.071, p=0.94). This allows us to rule out the number of motor responses as well as total earnings as sources of the difference in devaluation performance.

Neuroimaging results

Our primary objective was to investigate whether neural discrimination between the two instrumental conditions differed across acquisition and expression of behavioral control, particularly with respect to BOLD activity predictive of individual differences in devaluation insensitivity (see Table 2). To rule out the possibility that differences between our instrumental conditions were due to differences in the processing of relevant visual features (i.e., of abstract cues in the S-R condition and of the set of beakers in the R-O condition), additional analyses formally contrasted such differences with those emerging between matching control conditions (see Table 3). Unless otherwise noted, all effects reported below survived the subtraction of matching control conditions. All results described below survived correction for multiple comparisons at p<0.05 using either whole-brain cluster size thresholding (CST) or small volume corrections (SVC) based on coordinates from relevant previous studies. Notably, many of the areas specified a priori as targets of SVC, including the lingual gyrus, extra-striate cortex and inferior parietal lobule, also survived correction using whole-brain CST, as indicated in the fourth columns of Tables 1 and 2.

Discrimination between instrumental conditions during acquisition versus performance—

To delineate neural substrates engaged during early acquisition versus during expression of operant behavior, we entered the imaging data from the first two blocks of instrumental learning together with that from the subsequent test phase, for each instrumental condition, into a 2×2 (condition by experimental phase) analysis of variance (see Table 2). Although processes supporting acquisition should certainly dominate during the first two blocks of instrumental learning, we cannot completely rule out the possibility that some element of expression was also present during this phase. Inclusion of only the very earliest learning trials would not eliminate this possibility, but would introduce dramatically different error levels and severely reduce statistical power. We note, however, that the presence of processes related to expression during the early stages of instrumental learning should reduce, rather than enhance, discrimination between the levels of the "experimental phase" variable. Thus, in areas where we do observe a significant difference

in discriminatory activity across the different phases of the experiment, such differences are most likely attributable to the differential effects of acquisition versus expression.

S-R related activity: Activity that was greater in the S-R than the R-O condition (Fig. 2) emerged in the middle and superior occipital cortex, and in the superior parietal lobule (SPL). Interaction tests and (exclusively masked) tests of simple effects revealed that activity in the SPL and superior occipital cortex was significantly greater in the S-R than the R-O condition only during test performance, but not during early acquisition (Table 2). Moreover, during early acquisition, only a small area of the middle occipital cortex, the extra-striate body area (EBA; SVC), survived subtraction of corresponding matching controls (Table 3). In contrast, S-R selective activity throughout the occipital cortex, including the EBA, survived subtraction of matching controls during the test phase.

R-O related activity: Activity significantly greater in the R-O condition than the S-R condition (Fig. 3) was observed throughout a fronto-parietal network, including dorsal and ventral medial prefrontal cortex (VMPFC), the inferior parietal lobules (IPL), the posterior cingulate and the bilateral insula. Interaction tests and (exclusively masked) tests of simple effects revealed that, in the dorsomedial prefrontal cortex (DMPFC), discrimination between instrumental conditions occurred during test performance only (Table 2). Analyses assessing differences between conditions relative to matching controls also revealed R-O selective effects in the dorsal anterior caudate (aCN), but did not yield significant effects in the VMFPC (Table 3).

Discrimination between conditions across blocks of acquisition—To further assess differential activity related to acquisition processes, we tested for differences between conditions in training-dependent changes in neural activity, by adding linear weights to blocks of instrumental learning. An interaction test assessing neural activity that increased in the R-O condition and decreased in the S-R condition yielded significant effects throughout the right putamen and globus pallidus (CST, 33, -7, -5) and in the right IPL (CST, 63, -31, 37). The reverse interaction contrast, assessing neural activity that decreased in the R-O condition and increased in the S-R condition, did not reveal any significant effects. Specific assessments of increases versus decreases revealed *decreases* across training blocks in the S-R condition, but not in the (exclusively masked) R-O condition, throughout the right putamen and globus pallidus (Fig. 4 top). In contrast, significant *increases* in activity across blocks in the R-O condition, but not in the (exclusively masked) S-R condition, emerged in the right IPL (Fig. 4 bottom).

Neural correlates of devaluation insensitivity—To relate the neuroimaging data to our behavioral effects, we tested whether neural discrimination between instrumental conditions correlated with differences between conditions in the degree of devaluation insensitivity. This was indeed the case: Participants with stronger activation of the tail of caudate/thalamus (tCN/th) and of the cerebellum, extending into the lingual gyrus (LG), in the S-R relative to the R-O condition during the first two blocks of instrumental learning responded on a greater proportion of devalued trials in the S-R condition relative to the R-O condition during the subsequent test phase (Fig. 5A). These effects all survived exclusive

masking by an identical contrast applied, with a threshold of 0.1, to the imaging data from the test phase, suggesting that, in these areas, discriminatory activity during early acquisition, but *not* during test performance, predicted differences in devaluation insensitivity.

Conversely, during the test phase, differences in neural activity between the S-R and R-O conditions in the subgenual cortex (ROI) and ventral striatum (VS) were positively correlated with the difference between conditions in devaluation insensitivity (Fig. 5B). That is, *greater* subgenual and VS activity in the S-R relative to the R-O condition predicted greater devaluation insensitivity in the S-R than R-O condition. We also found a negative correlation with test performance in the right IPL (SVC), such that *lesser* IPL activity in the S-R relative to the R-O condition predicted greater devaluation insensitivity in the S-R than R-O condition. Again, the specificity of these results were confirmed using exclusive masking by an identical contrast, at a threshold of 0.1, applied to the imaging data from the training phase.

Discussion

In this study we explored the neural substrates of goal-directed and habitual action selection, with a focus on how recruitment of relevant brain areas might differ across acquisition and implementation of behavioral control. We scanned human participants with fMRI as they learned and performed a novel task designed to encourage either goal-directed encoding of the specific outcomes of instrumental responses (R-O condition), or a habitual mapping of responses to antecedent cues (S-R condition). In a subsequent test phase, participants were more likely to respond for a devalued outcome, indicative of habits, in the S-R condition. We found that neural activity in striatal and cortical areas 1) discriminated between our two instrumental conditions, 2), predicted individual differences in devaluation insensitivity and 3) did so differentially across acquisition and test performance.

Our results identified several areas in which the degree of neural discrimination between instrumental conditions predicted differences in devaluation insensitivity. Specifically, S-R selective activity in the tail of caudate and cerebellum during learning, but not during test performance, predicted greater devaluation insensitivity in the S-R, relative to the R-O, condition. The cerebellum has been strongly implicated in response automatization (Doyon et al., 1998; Doyon et al., 2002; Lang and Bastian, 2002; Balsters and Ramnani, 2011) – a resistance to dual-task interference postulated to be closely related to, and sometimes treated as synonymous with, habitual performance. Indeed, in the rodent literature, the cerebellum has been directly implicated in habit formation, such that cerebellar lesions abolish devaluation insensitivity in over-trained rats (Callu et al., 2007). However, to our knowledge, no previous study has directly linked the cerebellum to outcome devaluation insensitivity in humans. As with the cerebellum, the tCN has been implicated in skill learning (Poldrack and Gabrieli, 2001; Yamamoto et al., 2013). Moreover, single neuron recordings in the monkey tCN have revealed that, relative to the body and head of caudate, this area is specifically involved in encoding stable, non-flexible, values of visual cues (Kim and Hikosaka, 2013; Yamamoto et al., 2013; Hikosaka et al., 2014; Kim et al., 2014),

consistent with its role in devaluation insensitivity demonstrated here and elsewhere (Valentin et al., 2007; Liljeholm et al., 2012).

In the subgenual cingulate (BA 25), S-R selective activity during test performance, but not during learning, correlated with behavioral devaluation insensitivity. Based on its cytoarchitectonic subdivisions and connections, this area has been identified as homologous to the rodent infralimbic cortex: The rodent infralimbic and human subgenual areas both project heavily to the shell of the nucleus accumbence, and both are agranular, relatively poorly laminated, areas located ventrally on the medial wall (Gabbott et al., 1997; Ongur and Price, 2000; Ongur et al., 2003). In rodents, the infralimbic cortex has been shown to make specialized contributions to executive control processes facilitating the deployment of habits (Coutureau and Killcross, 2003; Haddon and Killcross, 2011; Smith et al., 2012; Smith and Graybiel, 2013). Consistent with such findings, our results suggest that a putative human homologue of this area does indeed play selective roles in the expression of habits. It should be noted, however, that the rodent infralimbic cortex is primarily known for its involvement in the extinction of Pavlovian responses (Milad and Quirk, 2002; Rhodes and Killcross, 2004; Rhodes and Killcross, 2007; Santini et al., 2008), while the human subgenual cortex has predominantly featured in studies on depression. (Greicius et al., 2007; Drevets et al., 2008; Johansen-Berg et al., 2008; Matthews et al., 2009; Keedwell et al., 2010). Future work is needed to reconcile these apparently divergent functions and their potential homology across species.

As with the subgenual cortex, activity in the ventral striatum (VS) during the test phase, but not during acquisition, predicted devaluation insensitivity. The VS has been frequently shown to support both the acquisition and expression of Pavlovian (stimulus-outcome) associations (Day et al., 2007; Blaiss and Janak, 2009), and to mediate the general motivational influence of such associations on instrumental performance; a phenomenon referred to as Pavlovian-instrumental transfer (Corbit and Balleine, 2011). Of particular importance for interpreting the current findings is the fact that, in rodents, PIT appears to selectively influence habitual, rather than goal-directed, responding: the greater the degree of insensitivity to outcome devaluation, the greater the general motivational influence of Pavlovian cues on instrumental performance (Holland, 2004; Balleine and Ostlund, 2007). We interpret the currently observed activity in the VS as reflecting a greater influence of Pavlovian cues on instrumental responding in the S-R than the R-O condition, and conjecture that this selective engagement of Pavlovian processes supported our behavioral effect.

Contrary to our predictions, we did not find a correlation between activity in the dorsal putamen during early learning and subsequent devaluation insensitivity. We did, however, find S-R selective decreases in right putamen activity across blocks of training – an effect that is consistent with a substantial literature demonstrating a decreased dependence on the putamen with extended (Ungerleider et al., 2002; Poldrack et al., 2005; Turner et al., 2005; Ashby et al., 2010), as well as intermediate (Brovelli et al., 2011), levels of training. Taken together, these results suggest that the contributions of the putamen to automatic and habitual behavioral control may take place early in acquisition, with long-term storage and mediation of well-trained performance occurring elsewhere (Orban et al., 2010). On the

other hand some neuroimaging studies have found increases in putamen activity with extended training (Floyer-Lea and Matthews, 2004; Tricomi et al., 2009; Wunderlich et al., 2012). For instance, Tricomi et al. (2009), found an increase in activity with overtraining in a far posterior region of the right putamen, concomitant with the development of habitual performance. One key difference between the Tricomi et al. study and the present one is that the S-R condition in the present study is optimized to generate rapid acquisition and expression of habitual behavior, whereas in the Tricomi et el. study, behavioral expression of habits (and perhaps also acquisition) emerged much more slowly, becoming evident in behavior only after several days of training. Thus, if the involvement of the putamen in S-R learning dissipates after a period of stable habitual performance, Tricomi et al. may not have sampled behavior beyond that stable period, whereas our accelerated habitual learning paradigm allowed us to do so. It is also worth noting, that whereas Tricomi et al. reported effects in a small area in the very far posterior putamen, our effects extend throughout the right putamen and globus pallidus.

One feature of the present results is that areas in which discriminatory neural activity was correlated with between-subject variation in devaluation insensitivity did not also show differential main effects in a comparison of S-R and R-O conditions. This may be in part due to the fact that our efforts to differentially encourage stimulus-response and response-outcome learning resulted only in relative differences in devaluation insensitivity, with the majority of participants showing some degree of sensitivity even in the S-R condition. Thus neural processes directly related to the degree of habitual responding might not be discernable in analyses that categorically compare conditions. On the other hand, not all aspects of encoding stimulus-response associations necessarily contribute to the dominance of such associations over performance. For example, the S-R selective decreases in neural activity across blocks of acquisition, found throughout the right putamen and globus pallidus, may reflect a gradual disengagement of processes that provide critical support during early development of stimulus-response associations – such as the maintenance of relevant representations, the encoding of discrepancies between attempted and accurate mappings, or the retrieval of response alternatives – but that are supplanted during subsequent performance by processes mediating behavioral control.

We also found selective recruitment by the S-R condition of a lateral middle occipital area that overlaps closely with what has been termed the “extrastriate body area” (EBA) (Downing et al., 2001). As the name implies, this region is known for its responses to images of human body parts, and has also been frequently implicated in action observation as well as execution (Astafiev et al., 2004; Kuhn et al., 2011; Liljeholm et al., 2012). Notably, studies assessing the role of the EBA in action execution have employed visually guided actions, such that an arbitrary visual stimulus indicates either the location towards which a movement should be directed (Astafiev et al., 2004), or which one of alternative actions should be performed (Kuhn et al., 2011). It is possible, therefore, that this area is specifically involved in limb movements that are largely stimulus-driven. In the current study, the S-R condition was designed to encourage a mapping of responses to arbitrary antecedent visual cues, by decorrelating actions from sensory-specific outcome features (and indeed from any visual features that were intrinsically related to the goal of maintaining beaker fluids). Thus, one possible explanation for the selective recruitment of EBA by the S-

R condition is that this area mediates the detection and mapping of arbitrary stimuli to behavioral responses.

The VMPFC has been implicated in goal-directed performance in numerous human neuroimaging studies, and is a proposed homolog of the rodent prelimbic cortex (Balleine and O'Doherty, 2010), lesions of which disrupt the acquisition, but not the expression, of goal-directed performance (Ostlund and Balleine, 2005). Although we did find selective recruitment of the VMPFC by the R-O condition during learning, this effect did not survive the subtraction of matching control conditions. It is possible, of course, that the process supported by the VMPFC, for example assigning values to subgoals, was elicited by the relevant stimuli in the matching task as well as the instrumental task. Another possibility, given our robust effect of experimental phase in the VMPFC, such that activity was greater during early acquisition than during test in both conditions, is that the goal-directed function supported by the VMPFC was present in both instrumental conditions during early acquisition. Further work is needed to arbitrate between these possibilities.

Further dissociating the contributions of medial prefrontal areas, we found that activity in a more dorsal medial prefrontal region was greater in the R-O than the S-R condition during test performance, but not during early acquisition. This region, along with adjacent dorsal anterior cingulate, has been implicated in tasks in which the attainment of goals and rewards requires high levels of cognitive control, such as when monitoring for unfavorable outcomes, and during response conflict and decision uncertainty (Ridderinkhof et al., 2004b; Ridderinkhof et al., 2004a). The currently observed pattern of activity in this area may reflect a decrease in performance monitoring and outcome evaluation during habitual relative to goal-directed control.

In our novel task the stimulus materials were identical across S-R and R-O conditions, but the instrumental contingencies encouraged participants to attend to different visual features (cues vs. beakers). Although we used a simple match-to-sample task to rule out visual processes involved in selectively attending to such features as a source of any imaging effects, there may be aspects of visual attention that are intrinsically related to instrumental responding. For example, we found greater activity in the R-O than the S-R condition in a posterior ventral region of the IPL (pvIPL) during both the learning and test phase; during test, discriminatory activity in this area predicted greater sensitivity to devaluation. The pvIPL has been proposed to function as a “circuit break” that re-directs attention towards behaviorally relevant information that is either exogenously presented (Corbetta and Shulman, 2002; Corbetta et al., 2008; Cabeza et al., 2012) or retrieved into working memory (Cabeza et al., 2012). Importantly, substantial evidence also indicates that the pvIPL is deactivated when the new information is not relevant to the current task (Shulman et al., 2003). As can be seen in Figure 3 (top), activity in the pvIPL appears to be deactivated in both instrumental conditions relevant to the matching control, but markedly more so in the S-R condition: A possible explanation for this pattern of results is that greater suppression was required in the S-R condition because the sensory features of the beaker subgoal, although ultimately irrelevant for response selection, was intimately related to the trial outcome. Indeed, this augmented suppression of sensory-specific outcome features may be a general property of S-R learning, particularly since the retrieval of such features into

working memory is considered central to goal-directed encoding (Balleine and Ostlund, 2007).

Consistent with evidence from the rodent literature that goal-directed performance persists in tasks in which alternative actions produce distinct rewards (Colwill and Rescorla, 1985; Holland, 2004), we found that human action selection was more goal-directed in a condition in which instrumental actions obtained unique sub-goals. Formally, the degree to which alternative actions yield distinct outcome states can be quantified as the *divergence* of their outcome probability distributions. In a previous study (Liljeholm et al., 2013) we found that activity in the right anterior IPL (the anterior dorsal supramarginal gyrus) increased with increasing outcome divergence. Likewise, in the current study, training-dependent increases in this area were found in the R-O, but not the S-R, condition, potentially reflecting the incremental increase in divergence, as actions became associated with their respective outcome states. The selective recruitment of the IPL by the R-O condition, and the correlation between such discriminatory neural activity and behavioral sensitivity to outcome devaluation, is also consistent with a large body of research implicating this area in various goal-directed processes, including the computation of instrumental contingencies (Seo et al., 2009; Liljeholm et al., 2011), the attribution of intent (den Ouden et al., 2005), awareness of agency (Farrer et al., 2008), and the mediation of executive function (Friedman and Goldman-Rakic, 1994). Further determination of the exact contributions of the IPL to goal-directed action selection, its role in corollary attentional processes, and in representations of outcome divergence, is an important avenue for future work.

In summary, our results are novel in several critical ways: To our knowledge, we are the first to dissociate the roles of the tail of caudate and ventral striatum, across learning and test performance, in behavioral insensitivity to outcome devaluation. We are also the first to demonstrate that activity in the inferior parietal lobule, an area that has been previously implicated in several processes closely linked to goal-directed action selection, including the attribution of intent and awareness of agency, predicts sensitivity to outcome devaluation. Finally, we reveal a potential functional homology between the human subgenual and rodent infralimbic cortex in the implementation of habitual control. Taken together, our findings suggest a broad systems division, at the cortical and subcortical levels, between brain areas mediating the acquisition and expression of action-outcome and stimulus-response associations. Notably, a fundamental issue in the search for treatments of behavioral disorders is how to both facilitate the automatization of actions that lead to healthful consequences and abolish well-established deleterious habits. An improved understanding of how distinct neural substrates in the human brain mediate the acquisition versus expression of habitual control, the aim of the present study, is therefore of significant clinical interest.

Acknowledgments

This work was supported by a grant from the Wellcome Trust and by an NIH grant (DA033077-01) to J.P.O.

References

Adams CD. Variations in the sensitivity of instrumental responding to refinforcer devaluation. Quarterly Journal of Experimental Psycholgoy. 1982; 34:77–98.

Adams CD, Dickinson A. Instrumental responding following reinforcer devaluation. Quarterly Journal of Experimental Psycholgoy. 1981; 33:109–121.

Ashby FG, Turner BO, Horvitz JC. Cortical and basal ganglia contributions to habit learning and automaticity. Trends Cogn Sci. 2010; 14:208–215. [PubMed: 20207189]

Astafiev SV, Stanley CM, Shulman GL, Corbetta M. Extrastriate body area in human occipital cortex responds to the performance of motor actions. Nat Neurosci. 2004; 7:542–548. [PubMed: 15107859]

Balleine BW, Dickinson A. Goal-directed instrumental action: contingency and incentive learning and their cortical substrates. Neuropharmacology. 1998; 37:407–419. [PubMed: 9704982]

Balleine BW, Ostlund SB. Still at the choice-point: action selection and initiation in instrumental conditioning. Ann N Y Acad Sci. 2007; 1104:147–171. [PubMed: 17360797]

Balleine BW, O'Doherty JP. Human and rodent homologies in action control: corticostriatal determinants of goal-directed and habitual action. Neuropsychopharmacology. 2010; 35:48–69. [PubMed: 19776734]

Balsters JH, Ramnani N. Cerebellar plasticity and the automation of first-order rules. J Neurosci. 2011; 31:2305–2312. [PubMed: 21307266]

Blaiss CA, Janak PH. The nucleus accumbens core and shell are critical for the expression, but not the consolidation, of Pavlovian conditioned approach. Behav Brain Res. 2009; 200:22–32. [PubMed: 19159648]

Brovelli A, Nazarian B, Meunier M, Boussaoud D. Differential roles of caudate nucleus and putamen during instrumental learning. Neuroimage. 2011; 57:1580–1590. [PubMed: 21664278]

Cabeza R, Ciaramelli E, Moscovitch M. Cognitive contributions of the ventral parietal cortex: an integrative theoretical account. Trends Cogn Sci. 2012; 16:338–352. [PubMed: 22609315]

Callu D, Puget S, Faure A, Guegan M, El Massioui N. Habit learning dissociation in rats with lesions to the vermis and the interpositus of the cerebellum. Neurobiology of disease. 2007; 27:228–237. [PubMed: 17560113]

Carelli RM, West MO. Representation of the body by single neurons in the dorsolateral striatum of the awake, unrestrained rat. J Comp Neurol. 1991; 309:231–249. [PubMed: 1885787]

Carelli RM, Wolske M, West MO. Loss of lever press-related firing of rat striatal forelimb neurons after repeated sessions in a lever pressing task. J Neurosci. 1997; 17:1804–1814. [PubMed: 9030639]

Colwill RM, Rescorla RA. Instrumental responding remains insensitive to reinfrocer devaluation after extensive training. J Exp Psychol Anim Behav Process. 1985; 11:520–536.

Corbetta M, Shulman GL. Control of goal-directed and stimulus-driven attention in the brain. Nature reviews Neuroscience. 2002; 3:201–215.

Corbetta M, Patel G, Shulman GL. The reorienting system of the human brain: from environment to theory of mind. Neuron. 2008; 58:306–324. [PubMed: 18466742]

Corbit LH, Balleine BW. The general and outcome-specific forms of Pavlovian-instrumental transfer are differentially mediated by the nucleus accumbens core and shell. J Neurosci. 2011; 31:11786–11794. [PubMed: 21849539]

Coutureau E, Killcross S. Inactivation of the infralimbic prefrontal cortex reinstates goal-directed responding in overtrained rats. Behav Brain Res. 2003; 146:167–174. [PubMed: 14643469]

Daw ND, Niv Y, Dayan P. Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. Nat Neurosci. 2005; 8:1704–1711. [PubMed: 16286932]

Day JJ, Roitman MF, Wightman RM, Carelli RM. Associative learning mediates dynamic shifts in dopamine signaling in the nucleus accumbens. Nat Neurosci. 2007; 10:1020–1028. [PubMed: 17603481]

de Wit S, Watson P, Harsay HA, Cohen MX, van de Vijver I, Ridderinkhof KR. Corticostriatal connectivity underlies individual differences in the balance between habitual and goal-directed action control. J Neurosci. 2012; 32:12066–12075. [PubMed: 22933790]

den Ouden HM, Frith U, Frith C, SJ B. Thinking about intentions. Neuroimage. 2005; 28:787–796. [PubMed: 15964210]

Dickinson A, Nicholas DJ, Adams CD. The effect of the instrumental training contingency on susceptibility to reinforcer devaluation. Quarterly Journal of Experimental Psycholgoy B. 1983; 35:35–51.

Downing PE, Jiang Y, Shuman M, Kanwisher N. A cortical area selective for visual processing of the human body. Science. 2001; 293:2470–2473. [PubMed: 11577239]

Doya K, Samejima K, Katagiri K, Kawato M. Multiple model-based reinforcement learning. Neural Comput. 2002; 14:1347–1369. [PubMed: 12020450]

Doyon J, Song AW, Karni A, Lalonde F, Adams MM, Ungerleider LG. Experience-dependent changes in cerebellar contributions to motor sequence learning. Proceedings of the National Academy of Sciences of the United States of America. 2002; 99:1017–1022. [PubMed: 11805340]

Doyon J, Laforce R Jr, Bouchard G, Gaudreau D, Roy J, Poirier M, Bedard PJ, Bedard F, Bouchard JP. Role of the striatum, cerebellum and frontal lobes in the automatization of a repeated visuomotor sequence of movements. Neuropsychologia. 1998; 36:625–641. [PubMed: 9723934]

Draganski B, Kherif F, Kloppel S, Cook PA, Alexander DC, Parker GJ, Deichmann R, Ashburner J, Frackowiak RS. Evidence for segregated and integrative connectivity patterns in the human Basal Ganglia. J Neurosci. 2008; 28:7143–7152. [PubMed: 18614684]

Drevets WC, Savitz J, Trimble M. The subgenual anterior cingulate cortex in mood disorders. CNS Spectr. 2008; 13:663–681. [PubMed: 18704022]

Esterman M, Tamber-Rosenau BJ, Chiu YC, Yantis S. Avoiding non-independence in fMRI data analysis: leave one subject out. Neuroimage. 2010; 50:572–576. [PubMed: 20006712]

Farrer C, Frey SH, Van Horn JD, Tunik E, Turk D, Inati S, Grafton ST. The angular gyrus computes action awareness representations. Cereb Cortex. 2008; 18:254–261. [PubMed: 17490989]

Floyer-Lea A, Matthews PM. Changing brain networks for visuomotor control with increased movement automaticity. J Neurophysiol. 2004; 92:2405–2412. [PubMed: 15381748]

Forman SD, Cohen JD, Fitzgerald M, Eddy WF, Mintun MA, Noll DC. Improved assessment of significant activation in functional magnetic resonance imaging (fMRI): use of a cluster-size threshold. Magn Reson Med. 1995; 33:636–647. [PubMed: 7596267]

Friedman HR, Goldman-Rakic PS. Coactivation of prefrontal cortex and inferior parietal cortex in working memory tasks revealed by 2DG functional mapping in the rhesus monkey. J Neurosci. 1994; 14:2775–2788. [PubMed: 8182439]

Gabbott PL, Dickie BG, Vaid RR, Headlam AJ, Bacon SJ. Local-circuit neurones in the medial prefrontal cortex (areas 25, 32 and 24b) in the rat: morphology and quantitative distribution. J Comp Neurol. 1997; 377:465–499. [PubMed: 9007187]

Glascher J. Visualization of group inference data in functional neuroimaging. Neuroinformatics. 2009; 7:73–82. [PubMed: 19140033]

Greicius MD, Flores BH, Menon V, Glover GH, Solvason HB, Kenna H, Reiss AL, Schatzberg AF. Resting-state functional connectivity in major depression: abnormally increased contributions from subgenual cingulate cortex and thalamus. Biological psychiatry. 2007; 62:429–437. [PubMed: 17210143]

Haddon JE, Killcross S. Inactivation of the infralimbic prefrontal cortex in rats reduces the influence of inappropriate habitual responding in a response-conflict task. Neuroscience. 2011; 199:205–212. [PubMed: 22015928]

Hampton AN, Bossaerts P, O'Doherty JP. The role of the ventromedial prefrontal cortex in abstract state-based inference during decision making in humans. J Neurosci. 2006; 26:8360–8367. [PubMed: 16899731]

Hikosaka O, Kim HF, Yasuda M, Yamamoto S. Basal ganglia circuits for reward value-guided behavior. Annual review of neuroscience. 2014; 37:289–306.

Holland PC. Relations between Pavlovian-instrumental transfer and reinforcer devaluation. J Exp Psychol Anim Behav Process. 2004; 30:104–117. [PubMed: 15078120]

Jog MS, Kubota Y, Connolly CI, Hillegaart V, Graybiel AM. Building neural representations of habits. Science. 1999; 286:1745–1749. [PubMed: 10576743]

Johansen-Berg H, Gutman DA, Behrens TE, Matthews PM, Rushworth MF, Katz E, Lozano AM, Mayberg HS. Anatomical connectivity of the subgenual cingulate region targeted with deep brain

stimulation for treatment-resistant depression. Cereb Cortex. 2008; 18:1374–1383. [PubMed: 17928332]

Keedwell PA, Drapier D, Surguladze S, Giampietro V, Brammer M, Phillips M. Subgenual cingulate and visual cortex responses to sad faces predict clinical outcome during antidepressant treatment for depression. Journal of affective disorders. 2010; 120:120–125. [PubMed: 19539998]

Killcross S, Coutureau E. Coordination of actions and habits in the medial prefrontal cortex of rats. Cereb Cortex. 2003; 13:400–408. [PubMed: 12631569]

Kim HF, Hikosaka O. Distinct basal ganglia circuits controlling behaviors guided by flexible and stable values. Neuron. 2013; 79:1001–1010. [PubMed: 23954031]

Kim HF, Ghazizadeh A, Hikosaka O. Separate groups of dopamine neurons innervate caudate head and tail encoding flexible and stable value memories. Frontiers in neuroanatomy. 2014; 8:120. [PubMed: 25400553]

Kuhn S, Keizer A, Rombouts SA, Hommel B. The functional and neural mechanism of action preparation: roles of EBA and FFA in voluntary action control. J Cogn Neurosci. 2011; 23:214–220. [PubMed: 20044885]

Lang CE, Bastian AJ. Cerebellar damage impairs automaticity of a recently practiced movement. J Neurophysiol. 2002; 87:1336–1347. [PubMed: 11877508]

Lee SW, Shimojo S, O'Doherty JP. Neural computations underlying arbitration between model-based and model-free learning. Neuron. 2014; 81:687–699. [PubMed: 24507199]

Liljeholm M, Molloy CJ, O'Doherty JP. Dissociable brain systems mediate vicarious learning of stimulus-response and action-outcome contingencies. J Neurosci. 2012; 32:9878–9886. [PubMed: 22815503]

Liljeholm M, Tricomi E, O'Doherty JP, Balleine BW. Neural correlates of instrumental contingency learning: differential effects of action-reward conjunction and disjunction. J Neurosci. 2011; 31:2474–2480. [PubMed: 21325514]

Liljeholm M, Wang S, Zhang J, O'Doherty JP. Neural correlates of the divergence of instrumental probability distributions. J Neurosci. 2013; 33:12519–12527. [PubMed: 23884955]

Maldjian JA, Laurienti PJ, Kraft RA, Burdette JH. An automated method for neuroanatomic and cytoarchitectonic atlas-based interrogation of fMRI data sets. Neuroimage. 2003; 19:1233–1239. [PubMed: 12880848]

Matthews S, Simmons A, Strigo I, Gianaros P, Yang T, Paulus M. Inhibition-related activity in subgenual cingulate is associated with symptom severity in major depression. Psychiatry research. 2009; 172:1–6. [PubMed: 19239982]

Milad MR, Quirk GJ. Neurons in medial prefrontal cortex signal memory for fear extinction. Nature. 2002; 420:70–74. [PubMed: 12422216]

O'Doherty JP, Deichmann R, Critchley HD, Dolan RJ. Neural responses during anticipation of a primary taste reward. Neuron. 2002; 33:815–826. [PubMed: 11879657]

Ongur D, Price JL. The organization of networks within the orbital and medial prefrontal cortex of rats, monkeys and humans. Cereb Cortex. 2000; 10:206–219. [PubMed: 10731217]

Ongur D, Ferry AT, Price JL. Architectonic subdivision of the human orbital and medial prefrontal cortex. J Comp Neurol. 2003; 460:425–449. [PubMed: 12692859]

Orban P, Peigneux P, Lungu O, Albouy G, Breton E, Laberenne F, Benali H, Maquet P, Doyon J. The multifaceted nature of the relationship between performance and brain activity in motor sequence learning. Neuroimage. 2010; 49:694–702. [PubMed: 19732838]

Ostlund SB, Balleine BW. Lesions of medial prefrontal cortex disrupt the acquisition but not the expression of goal-directed learning. J Neurosci. 2005; 25:7763–7770. [PubMed: 16120777]

Poldrack RA, Gabrieli JD. Characterizing the neural mechanisms of skill learning and repetition priming: evidence from mirror reading. Brain : a journal of neurology. 2001; 124:67–82. [PubMed: 11133788]

Poldrack RA, Sabb FW, Foerde K, Tom SM, Asarnow RF, Bookheimer SY, Knowlton BJ. The neural correlates of motor skill automaticity. J Neurosci. 2005; 25:5356–5364. [PubMed: 15930384]

Rhodes SE, Killcross S. Lesions of rat infralimbic cortex enhance recovery and reinstatement of an appetitive Pavlovian response. Learn Mem. 2004; 11:611–616. [PubMed: 15466316]

Rhodes SE, Killcross AS. Lesions of rat infralimbic cortex enhance renewal of extinguished appetitive Pavlovian responding. Eur J Neurosci. 2007; 25:2498–2503. [PubMed: 17445245]

Ridderinkhof KR, van den Wildenberg WP, Segalowitz SJ, Carter CS. Neurocognitive mechanisms of cognitive control: the role of prefrontal cortex in action selection, response inhibition, performance monitoring, and reward-based learning. Brain Cogn. 2004a; 56:129–140. [PubMed: 15518930]

Ridderinkhof KR, Ullsperger M, Crone EA, Nieuwenhuis S. The role of the medial frontal cortex in cognitive control. Science. 2004b; 306:443–447. [PubMed: 15486290]

Santini E, Quirk GJ, Porter JT. Fear conditioning and extinction differentially modify the intrinsic excitability of infralimbic neurons. J Neurosci. 2008; 28:4028–4036. [PubMed: 18400902]

Seo H, Barraclough DJ, Lee D. Lateral intraparietal cortex and reinforcement learning during a mixed-strategy game. J Neurosci. 2009; 29:7278–7289. [PubMed: 19494150]

Shulman GL, McAvoy MP, Cowan MC, Astafiev SV, Tansy AP, d'Avossa G, Corbetta M. Quantitative analysis of attention and detection signals during visual search. J Neurophysiol. 2003; 90:3384–3397. [PubMed: 12917383]

Smith KS, Graybiel AM. A dual operator view of habitual behavior reflecting cortical and striatal dynamics. Neuron. 2013; 79:361–374. [PubMed: 23810540]

Smith KS, Virkud A, Deisseroth K, Graybiel AM. Reversible online control of habitual behavior by optogenetic perturbation of medial prefrontal cortex. Proceedings of the National Academy of Sciences of the United States of America. 2012; 109:18932–18937. [PubMed: 23112197]

Tanaka SC, Balleine BW, O'Doherty JP. Calculating consequences: brain systems that encode the causal effects of actions. J Neurosci. 2008; 28:6750–6755. [PubMed: 18579749]

Tang C, Pawlak AP, Prokopenko V, West MO. Changes in activity of the striatum during formation of a motor habit. Eur J Neurosci. 2007; 25:1212–1227. [PubMed: 17331217]

Tricomi E, Balleine BW, O'Doherty JP. A specific role for posterior dorsolateral striatum in human habit learning. Eur J Neurosci. 2009; 29:2225–2232. [PubMed: 19490086]

Turner, RS.; McCairn, K.; Simmons, D.; Bar-Gad, I. Sequential motor behavior and the basal ganglia. In: Bolam, JP.; Ingham, CA.; Magill, PJ., editors. The basal ganglia VIII (Advances in Behavioral Biology). Springer; 2005. p. 563-574.

Ungerleider LG, Doyon J, Karni A. Imaging brain plasticity during motor skill learning. Neurobiol Learn Mem. 2002; 78:553–564. [PubMed: 12559834]

Valentin VV, Dickinson A, O'Doherty JP. Determining the neural substrates of goal-directed learning in the human brain. J Neurosci. 2007; 27:4019–4026. [PubMed: 17428979]

Wunderlich K, Dayan P, Dolan RJ. Mapping value based planning and extensively trained choice in the human brain. Nat Neurosci. 2012; 15:786–791. [PubMed: 22406551]

Yamamoto S, Kim HF, Hikosaka O. Reward value-contingent changes of visual responses in the primate caudate tail associated with a visuomotor skill. J Neurosci. 2013; 33:11227–11238. [PubMed: 23825426]

Yin HH, Knowlton BJ, Balleine BW. Lesions of dorsolateral striatum preserve outcome expectancy but disrupt habit formation in instrumental learning. Eur J Neurosci. 2004; 19:181–189. [PubMed: 14750976]

Yin HH, Knowlton BJ, Balleine BW. Blockade of NMDA receptors in the dorsomedial striatum prevents action-outcome learning in instrumental conditioning. Eur J Neurosci. 2005; 22:505–512. [PubMed: 16045503]

Abbreviations

aCN	anterior caudate nucleus
ANOVA	analysis of variance
BA	Brodmann area
BOLD	blood oxygenation level dependent

CRBL	Cerebellum
CST	cluster-size threshold
DLS	Dorsolateral striatum
DMPFC	Dorsomedial prefrontal cortex
EBA	Extrastriate body area
fMRI	functional magnetic resonance imaging
GLM	general linear model
IPL	inferior parietal lobule
ITI	inter-trial interval
LG	lingual gyrus
LOSO	leave one subject out
MNI	Montreal Neurological Institute
mOFC	medial orbitofrontal
ms	milliseconds
PIT	Pavlovian-instrumental transfer
Post-/Pre-central S.	Post-/Pre-central Sulcus
Post. Cing	Posterior Cingulate
pvIPL	posterior ventral IPL
R-O	response-outcome
ROI	region of interest
S-R	stimulus-response
SEM	standard error of the mean
SMA	supplementary motor area
SPL	superior parietal lobule
SPM	statistical parametric mapping
SVC	small volume correction
tCN	tail of caudate
th	thalamus
VLPFC	Ventrolateral prefrontal cortex
VMPFC	ventromedial prefrontal cortex
VS	ventral striatum
WMA	world medical association

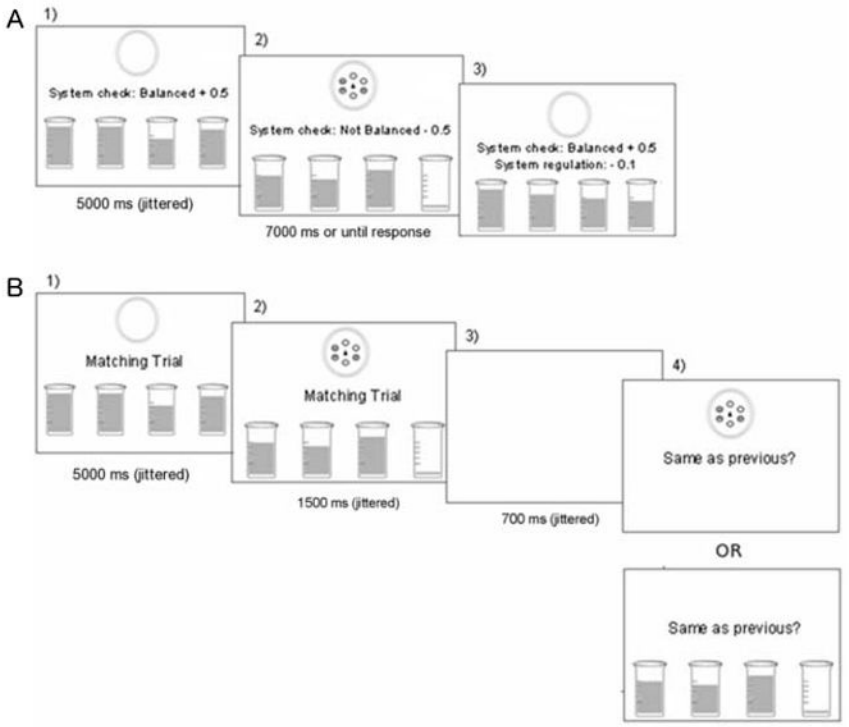

Figure 1.

Illustrations of instrumental learning and matching control tasks. A. Instrumental Learning: Participants are required to maintain the balance of a system of fluid-filled beakers using a set of four instrumental actions (each a 3-press sequence, see methods). During the inter-trial interval, the liquid in the beakers continually fluctuate but remain high, and "balance checks", occurring at brief random intervals, yield points for system balance (1). At the trial onset, one of four abstract cues appears, the liquid in one of the beakers drops to its bottom, and balance checks begin to indicate a loss of points due to system imbalance (2). Points are continually lost until the participant successfully re-fills the emptied beaker using one of the four actions. Following completion of the correct action (3), the abstract cue disappears, the beaker is re-filled, a small fee is charged for regulating the system, and balance checks again yield points for system balance. If the correct action is not performed within 7 seconds, the beaker is automatically re-filled, in which case there is no charge for system regulation. B. Matching task: The inter-trial interval (1) and trial onset (2) were as in the instrumental task but were followed, 1500 ms after trial onset, by a blank screen with a 700 ms duration (3). The subsequent, final, screen (4) showed a matching/non-matching stimulus together with a query about the match. In the S-R condition, the final screen always showed one of the abstract cues (top); conversely, in the R-O condition, the stimulus to-be-matched was always a set of beakers (bottom).

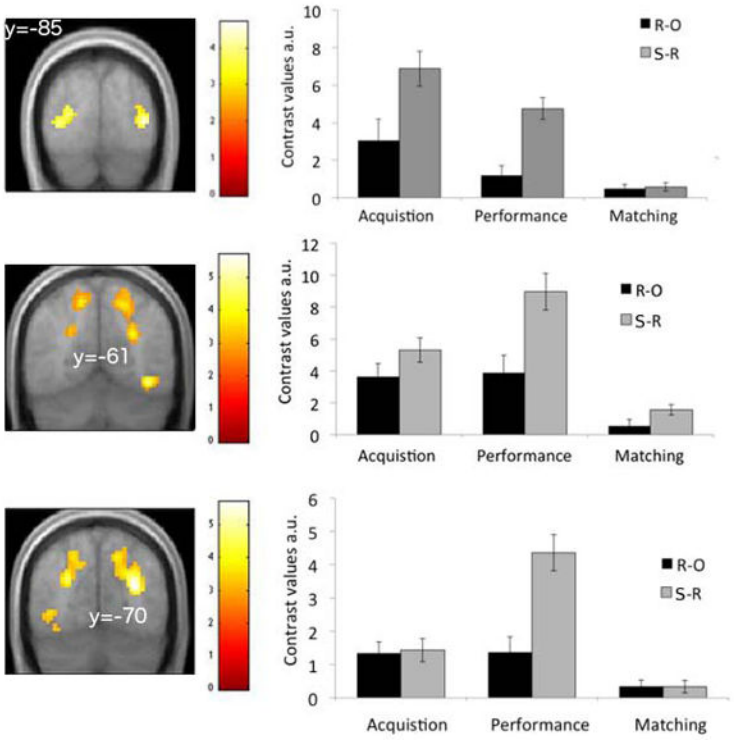

Figure 2.

Areas selectively involved in the Stimulus-Response condition. Statistical maps show main effects of the S-R > R-O contrast, in the middle occipital cortex (EBA) (top row), and an interaction effect (i.e., [S-R > R-O (performance > acquisition)]) in the superior occipital cortex, extending into the SPL (middle and bottom rows). Bar plots show contrast values estimated at LOSO coordinates for each instrumental condition, during both acquisition and performance, and for matching controls. Error bars = SEM.

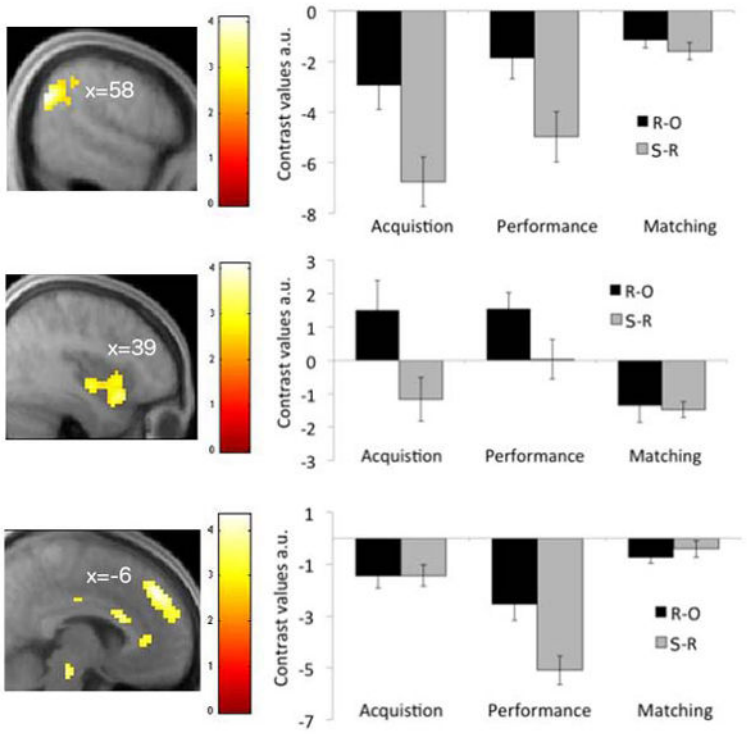

Figure 3.

Areas selectively involved in the Response-Outcome condition. Statistical maps show main effects of the R-O > S-R contrast in the posterior ventral inferior parietal lobule (top) and insula (middle), as well as an interaction effect (i.e., [R-O > S-R (performance > acquisition)]) in the DMPFC (bottom). Bar plots show contrast values estimated at LOSO coordinates for each instrumental condition, during both acquisition and performance, and for matching controls. Error bars = SEM.

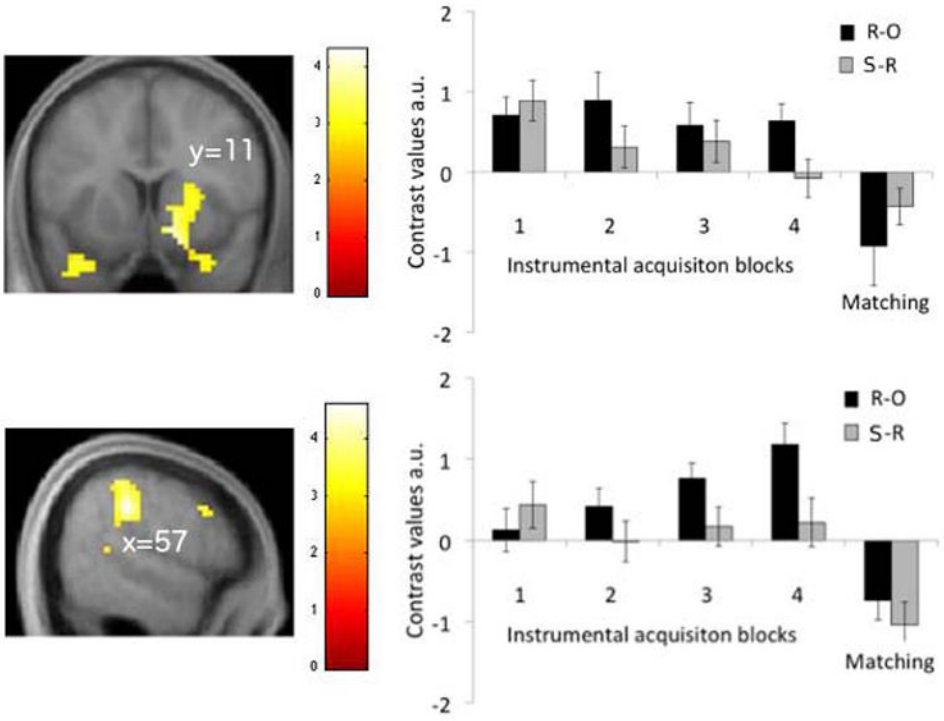

Figure 4.

Discrimination between conditions across blocks of acquisition. A. Statistical map showing results from disjunction tests assessing *decreases* across blocks of acquisition in the S-R but *not* the R-O condition with effects emerging in the right putamen/globus pallidus. B. Results from a disjunction test of *increases* across blocks in the R-O but *not* the S-R condition yielded effects in the IPL. Bar plots show contrast values estimated at LOSO coordinates for each instrumental condition in each training block, as well as for matching control conditions. Error bars = SEM.

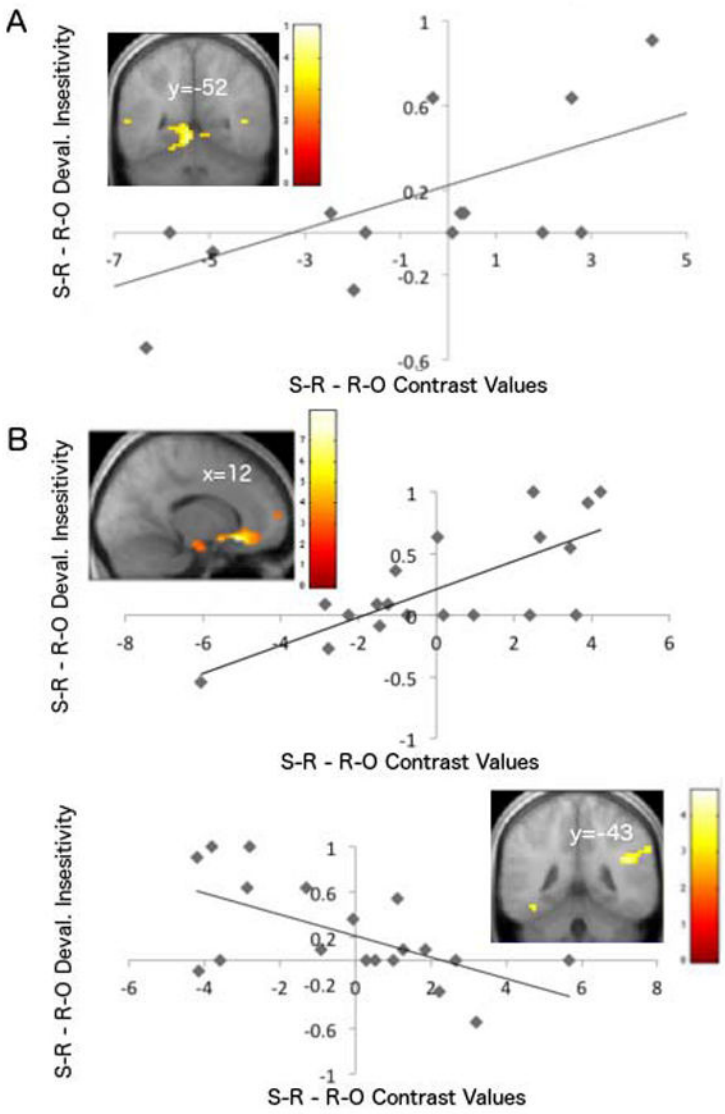

Figure 5.

Correlations between contrast values for the [S-R > R-O (instrumental > control)] contrast and behavioral differences between the S-R and R-O conditions in the proportion of devaluation insensitive responses. Contrast values are extracted from 8 mm spheres centered on the peak voxel in each area. A. Correlation between contrast values from the first two blocks of instrumental learning (x-axis) and differences between instrumental conditions in the proportion of devaluation insensitive responses (y-axis), showing effects in the cerebellum. B. Correlation between contrast values from the test phase (x-axis) and concurrent differences between instrumental conditions in the proportion of devaluation insensitive responses (y-axis), showing a positive correlation in the subgenual cortex (top) and a negative correlation in the right IPL (bottom).

Table 1

Imaging results from a 2×2 ANOVA contrasting instrumental conditions and training phases.

Test	Area	x, y, z	T at peak level	Correction	Cluster size at p<0.005
Main effects					
SR > RO	Inferior Occipital	36, -85, -8	6.33	CST	859
	Middle Occipital	36, -85, 7	6.82		
	Superior Occipital	27, -67, 28	5.16		
	SPL	18, -61, 52	4.22		
RO > SR	DMPFC	18, 56, 25	3.62	CST	843
	VMPFC	3, 41, 1	3.25	CST	
	IPL	60, -55, 28	3.99	CST	187
	Posterior Cing.	-3, -34, 43	3.13	CST	190
	Insula	39, 14, -41	3.89	CST	157
Acq. > Perf.	VMPFC	-3, 62, 4	6.12	CST	2181
	DMPFC	-3, 47, 37	5.65		
	VS	-6, 5, -8	4.90		
	aCN	9, 20, 7	3.53		
	VLPFC	-45, 32, -11	5.46	CST	354
	IPL	-51, -64, 34	4.76	CST	200
Perf. > Acq.	Cerebellum	27, -61, -26	9.17	CST	9773
	SMA	0, 8, 46	5.15		
	Insula/Putamen	-33, 8, 4	5.69		
	Postcentral S.	-42, -37, 58	6.33		
	Precentral S.	-30, -4, 52	5.97		
	Calcarine	24, -58,10	4.94		
	SPL	-15, -64, 49	5.31		
Interactions					
SR > RO (Perf.>Acq.)	Occipital	33, -79, 19	5.10	CST	859
	SPL	24, -52, 46	3.58		
RO > SR (Perf.>Acq.)	DMPFC	-9, -41, 40	4.26	CST	168

* The absence of an entry in the "correction" and "cluster size" columns indicates that the relevant cluster is continuous (at the $p<0.005$ threshold) with that listed above it. SPL=Superior Parietal Lobule, V/DMPFC=Ventro-/Dorsomedial Prefrontal Cortex, IPL=Inferior Parietal Lobule, aCN=anterior caudate nucleus, Post. Cing.=Posterior Cingulate, VS=Ventral Striatum, VLPFC=Ventrolateral Prefrontal Cortex, SMA=Supplementary Motor Area, Post-/Pre-central S.= Post-/Pre-central Sulcus.

Table 2

Imaging results from 2×2 ANOVAs contrasting instrumental and matching tasks, and from tests of neural correlates of devaluation insensitivity

Contrast	Area	x, y, z	T at peak level	* Correction	* Cluster size at p<0.005
Learning phase					
SR > RO (instr > ctrl)	EBA	39, -79, 10	4.09	SVC	45
RO > SR (instr > ctrl)	IPL	57, -58, 31	4.41	CST	275
	Insula	33, 14, 4	3.57	CST	421
	Precuneus	6, -70, 37	4.28	CST	395
	Post. Cing.	9, -31, 43	3.36		
	aCN	12, 2, 16	3.40	SVC	18
Test phase					
SR > RO (instr > ctrl)	Occipital	30, -79, 19	5.30	CST	1153
	SPL	24, -62, 49	4.47		
RO > SR (instr > ctrl)	IPL	57, -46, 37	3.09	SVC	18
	Insula	42, -1, -5	4.06	CST	383
	DMPFC	-6, 41, 40	3.87	CST	382
Matching effects					
SR control > RO control	Lateral LG	24, -91, -8	5.20	CST	198
RO control > SR control	Medial LG	-9, -76, -8	5.15	CST	1541
	Cuneus	-9, -91, 19	5.01		
Correlates of devaluation insensitivity					
Positive correlation	CRBL/LG	-6, -52, -8	5.06	CST	125
	tCN/th	-12, -28, 10	4.19	SVC	21
	Subgenual	-15, 23, -14	5.24	CST	224
	VS	-9, 5, -14	8.47		
Negative correlation	IPL	45, -43, 34	4.62	SVC	54

* The absence of an entry in the “correction” and “cluster size” columns indicates that the relevant cluster is continuous (at the p<0.005 threshold) with that listed above it. EBA=Extrastriate Body Area, IPL=Inferior Parietal Lobule, Post. Cing.=Posterior Cingulate aCN=Anterior caudate nucleus, SPL=Superior Parietal Lobule, DMPFC= Dorsomedial Prefrontal Cortex, VS=Ventral Striatum, tCN/th=tail of caudate/thalamus, CRBL=Cerebellum, LG=Lingual Gyrus.