Neural mechanisms of cue-approach training

Akram Bakkour[a,b,1], **Jarrod A. Lewis-Peacock**[a,b,c], **Russell A. Poldrack**[a,b,c,2], and **Tom Schonberg**[a,d,*]

[a]Imaging Research Center, The University of Texas at Austin, 100 E 24th St, Stop R9975, Austin, TX, 78712, USA

[b]Department of Neuroscience, The University of Texas at Austin, 100 E 24th St, Stop C7000, Austin, TX, 78712, USA

[c]Department of Psychology, The University of Texas at Austin, 108 E Dean Keeton, Stop A8000, Austin, TX, 78712, USA

[d]Department of Neurobiology, Faculty of Life Sciences and Sagol School of Neuroscience, Tel Aviv University, P.O. Box 39040, Tel Aviv, 6997801, Israel

Abstract

Biasing choices may prove a useful way to implement behavior change. Previous work has shown that a simple training task (the cue-approach task), which does not rely on external reinforcement, can robustly influence choice behavior by biasing choice toward items that were targeted during training. In the current study, we replicate previous behavioral findings and explore the neural mechanisms underlying the shift in preferences following cue-approach training. Given recent successes in the development and application of machine learning techniques to task-based fMRI data, which have advanced understanding of the neural substrates of cognition, we sought to leverage the power of these techniques to better understand neural changes during cue-approach training that subsequently led to a shift in choice behavior. Contrary to our expectations, we found that machine learning techniques applied to fMRI data during non-reinforced training were unsuccessful in elucidating the neural mechanism underlying the behavioral effect. However, univariate analyses during training revealed that the relationship between BOLD and choices for Go items increases as training progresses compared to choices of NoGo items primarily in lateral prefrontal cortical areas. This new imaging finding suggests that preferences are shifted via differential engagement of task control networks that interact with value networks during cue-approach training.

[*]Correspondence: Tom Schonberg, Department of Neurobiology, Faculty of Life Sciences and Sagol School of Neuroscience, Tel Aviv University, P.O. Box 39040, Tel Aviv, 6997801, Israel. schonberg@post.tau.ac.il.
[1]**Present address:** Department of Psychology, Columbia University, 1190 Amsterdam Ave, MC5501, New York, NY, 10027, USA
[2]**Present address:** Department of Psychology, Stanford University, Bldg. 420, Jordan Hall, Stanford, CA, 94305, USA

Keywords

cue-approach training; behavioral change; value-based decision making

1. Introduction

In order to eliminate unhealthy behaviors, one must find ways to enhance healthy choices. Changing preferences is an important strategy in addressing public health concerns, such as the obesity epidemic. To achieve lasting behavioral change to improve health, one must overcome the automaticity and strength of first-learned habits. First-learned behaviors are the rule that must be broken by subsequent learning in order for new habits to replace older ones over the long term (Bouton, 2004). Initial positive change in behavior may be achieved through intervention based on willful effort (Schonberg et al., 2014b; Tricomi et al., 2009), but the long term prospects for such improvement are uncertain (Bjork, 2001; Bouton, 1993; Cahill and Perera, 2011; Higgins et al., 1995; Wood and Neal, 2007). Focus has turned to targeting automatic processes to change human behavior with the goal of preventing disease (Marteau et al., 2012).

Previous research on value-based decision making has focused mostly on external reinforcement (O'Doherty et al., 2004; Thorndike, 1911) or the description of the decision problem (De Martino et al., 2006; Slovic, 1995; Tversky and Kahneman, 1986), but few have attempted to directly influence the underlying subjective values of individual options. In previous work by our group, we showed that choices can be biased toward targeted food items and the subjective value placed on these items can be differentially modulated by simply associating particular food items with an auditory cue to perform a motor response, without relying on external reinforcement or reframing the decision problem (Schonberg et al., 2014a). The previously described cue-approach task (CAT) is similar to the cued inhibition version of the stop-signal task (Lenartowicz et al., 2011; Verbruggen and Logan, 2008), with a crucial difference. In a typical stop-signal task, participants press a button on the keyboard every time a stimulus appears on the screen, except when a tone sounds they must try to inhibit a prepotent motor response. In CAT however, participants passively view stimuli on the screen, except when a tone sounds, they must press a button on the keyboard as quickly as possible. Training inhibition has been demonstrated to influence choice behavior for appetitive stimuli (Houben et al., 2012; Lawrence et al., 2015; Veling et al., 2013) and value for neutral stimuli (Wessel et al., 2014). Following stop-signal or go/no-go inhibition, participants tended to avoid or devalue stimuli that were associated with inhibition of action. However, rather than aiming to decrease choices, we developed CAT seeking to enhance choices for certain stimuli. In the original version of CAT, participants were asked to fast for four hours prior to arriving for the experiment. After providing informed consent, they were endowed with $3 to take part in an auction to obtain their pre-experimental preferences for 60 food items (Becker et al., 1964; Plassmann et al., 2007). Items were then rank ordered based on preference and median split into high and low value items. High and low value items were then placed into one of two experimental conditions: Go or NoGo. During training, participants passively viewed pictures of food items and pressed a button when they heard an infrequent tone. In a subsequent probe phase,

participants chose one item from a pair of equally preferred items, one associated with a tone during training (Go) and the other not associated with a tone (NoGo). Cue-approach training has proven to directly influence preference for single items through choice behavior following training. Approached (Go) items were chosen more often than initially equally preferred, non-approached (NoGo) items (Bakkour et al., 2016; Schonberg et al., 2014a). This procedure successfully changed choice behavior and the effect was maintained over six to eight weeks for participants who underwent the longest training (Schonberg et al., 2014a). Such a shift in choice behavior is thought to be mediated by an increase in gain in the coding of value for Go items in the ventromedial prefrontal cortex (vmPFC, Schonberg et al., 2014a), a brain region that has previously been heavily implicated in coding for value (Bartra et al., 2013; Padoa-Schioppa and Assad, 2006). This work has established cue-approach training as a model for non-reinforced preference change via modulation of subjective value for individual items. The question remains; how are values of Go items being modulated during CAT training?

Development of CAT was influenced by work on the attentional boost effect (Lin et al., 2010; Swallow and Jiang, 2010). In a typical attentional boost task, participants have better subsequent memory for incidental stimuli that were presented along with targets than those that were presented along with non-targets. The attentional boost effect established the importance of behavioral relevance in improving memory for incidental information. The cue-approach effect similarly established the importance of behavioral relevance for shifting preferences. Follow-up behavioral studies (Bakkour et al., 2016), using variations on the basic cue-approach training task have singled out memory retrieval and sustained top-down attention mechanisms to be at play during cue-approach training, leading to a shift in preferences at a later choice phase. However, standard univariate analyses of training-phase fMRI data in the previous imaging study of CAT were inconclusive and did not provide any insight into the neural mechanism responsible for modulating values of individual items during CAT training (Schonberg et al., 2014a). In the current study, we set out to characterize changes in neural activity during the cue-approach training phase using both univariate and multivariate analysis techniques.

Machine learning and pattern recognition algorithms have recently been adapted and developed to decode and characterize cognitive task-relevant neural activity using fMRI data (see Lemm et al., 2011; Mahmoudi et al., 2012, for review). One of the most popular of these machine-learning techniques is linear classification. This is a technique for decoding information about task variables from patterns of activity across an array of voxels. One of the common linear classification algorithms is the linear support vector machine (SVM). In this study, we sought to train a linear SVM classifier to identify whole-brain fMRI patterns elicited by cognitive processes thought to underlie shifts in choice preference during cue-approach training. Our hypothesis was that changes in classifier identification of the level of engagement of these cognitive processes of interest during training would predict later choices, reflecting a shift in preferences.

In order to test our hypothesis, we developed a cognitive localizer task that engages three distinct cognitive processes implicated in value change during the cue-approach training task: perception, memory retrieval, and valuation. We used multivariate pattern analysis

techniques on fMRI data acquired during this novel task to predict the level of engagement of these cognitive processes during cue-approach training. We investigated how changes in these processes (as measured by classifier predictions) contributed to a shift in preferences at a later choice phase. This analysis allows us to directly test our hypothesis that changes in the level of engagement of these particular cognitive processes during training predicts a shift in choice behavior. We expect that increases in the engagement of valuation and memory retrieval processes over the course of cue-approach training will be related to later choices. Furthermore, we were able to test whether process engagement progressed differentially for Go and NoGo trials as training proceeded. We predicted that the differential change in engagement of valuation and memory retrieval processes from beginning to end of the training phase, rather than the difference in overall engagement of these processes, would be predictive of later choices as participants learn to associate the food item with the tone cue as the training phase progresses. This allowed us the potential to better understand the neural mechanisms underlying non-reinforced training that leads to a shift in preferences. Finally, we also used standard univariate fMRI analysis techniques on probe phase data to replicate previous findings, and on training phase data to identify changes in whole-brain activation throughout training.

The design of the current study was optimized for application of MVPA techniques to identify underlying neurocognitive mechanisms for the CAT effect. Previous studies have demonstrated the power of these techniques not only to classify distributed patterns of fMRI activity elicited by different categories of images while the participant was viewing them (Cox and Savoy, 2003; Haxby et al., 2001), but also to classify intentions (Haynes et al., 2007; Soon et al., 2008), attentional states (Rosenberg et al., 2015) and the contents of memory recall (Polyn et al., 2005) using classifiers trained on different sets of stimuli from those being classified. Furthermore, and most germane to our main question of interest in the current study, Gross et al. (2014) trained an SVM classifier to discriminate levels of subjective value of foods and predicted the subjective value of engaging activities and vice versa. This supports the idea of common representation of value and the valuation process across domains. This finding also suggests that classifying the valuation process in one task can be used to decode value from a different task as planned in the current study. Other studies demonstrated robust cross-modal or cross-task classification (Lewis-Peacock et al., 2012; 2015). Polyn et al. (2005) trained classifiers on fMRI data from a localizer task requiring the perception and evaluation of familiar pictures, and then used these classifiers to decode the category of stimuli being retrieved from long-term memory during free recall. Lewis-Peacock and Postle (2008) used the same localizer task and analysis approach to decode the contents of working memory during cued recall. Esterman and colleagues (2009) used fMRI pattern classifiers to decode which domain of cognitive control (e.g., shifting visuospatial attention, switching task rules, shifting attention in working memory) was engaged at any given moment. Together, these findings suggest that fMRI classifiers trained on long-term memory retrieval might be able to identify the engagement of this process during CAT training.

2. Materials and Methods

2.1. Participants

Thirty-two healthy right-handed participants (17 female, mean age = 21.8 ± 3.1, age range: 18-29, mean body mass index (BMI) = 22.3 ± 3.8) completed the standard CAT while in a magnetic resonance imaging (MRI) scanner.

All participants had normal or corrected-to-normal vision, no history of psychiatric, neurologic or metabolic illnesses, no history of eating disorders, no food restrictions, and were not taking any medications that would interfere with the experiment. Participants were also free of any metal implants or any other contraindications for MRI. Participants were told that the goal of the experiment was to study food preferences and were asked to refrain from eating for four hours prior to arrival at the laboratory (Plassmann et al., 2007). All participants gave informed consent and the institutional review board (IRB) at the University of Texas at Austin approved the study.

2.2. Task

2.2.1. Auction—After consenting to take part in the study and filling out standard MRI safety metal screening forms, participants were endowed with $3, which they used to take part in an auction. Participants were presented with one snack item at a time on a computer screen. Food items were presented in random order. They placed their bid by moving the mouse cursor along an analog scale that spanned from 0 to 3 at the bottom of the screen. The auction was self-paced and the next item was presented only after the participants placed their bid. The auction procedure allowed us to obtain a measure of willingness to pay (WTP) for each of 56 appetitive food items per participant. The auction followed the BDM rules (Becker et al., 1964). Participants were told that their best strategy to win the auction was to bid exactly what each item was worth to them to purchase from the experimenter at the end of the experiment and that bidding consistently high or consistently low was a bad strategy. They were told that a single trial would be drawn at random at the end of the session and that they could use any amount of the full $3 for each food item and would not be spreading their endowment over multiple items. At the end of the session, the computer generated a counter bid; a random number between $0 and $3 in 25 cent increments. If the computer bid was equal to or higher than the participant's bid, then he or she lost the auction. If, however, the participant outbid the computer, then they were offered to purchase the randomly drawn food item from the experimenter at the computer's bid lower price.

2.2.2. Item Selection—Items were ranked based on WTP, where item #1 had the highest WTP and item #56 the lowest. 24 items with fixed rank order numbers from the full range were selected to serve as stimuli for the cognitive localizer task (see section 2.2.3 below). From the remaining 32 items, eight items were designated as higher-valued (from items ranked 8 through 18) and eight items as lower-valued (from items ranked 39 though 49). Out of each of these eight items, four were associated with an auditory cue (Go items) and four without any cue to press a button (NoGo items) during the cue-approach task (Figure 1). This selection procedure ensured pairing of high-value Go with high-value NoGo items and low-value Go with low-value NoGo items such that items in each pair later presented at

probe were on average matched for WTP. Based on initially stated values during the auction, participants should a priori be indifferent in choosing between items in these pairs. Of the 56 items presented during auction, 24 were used during the cognitive localizer task and the other 32 items were used during training. To maintain 25% cue frequency as is standard in stop-signal tasks (Logan and Cowan, 1984), 24 out of 32 items were NoGo items during training. Eight Go and eight NoGo items were presented in pairs during probe. Of those 32 items that were used during training, only 16 were used during the probe phase (Figure 1B). Item assignment to Go and NoGo conditions to be used later in the probe phase was counterbalanced across participants.

2.2.3. Cognitive Localizer—In this task, participants were presented with one food stimulus at a time and at the bottom of the screen one of three questions appeared (Figure 2A). Participants were asked to answer the question relevant to the item on the screen. Each of 24 items appeared with all three questions in random order over two runs. The three questions require three distinct cognitive processes: 1) Valuation: "How much would you like to eat this item?" Four alternative forced choices were ranked from 1 (most) to 4 (least). 2) Memory retrieval: "When did you last see this item at a store?" Four alternative forced choices from never to within the last week. 3) Perceptual decision: "How many items are outside the packaging?" Two alternative forced choices, either one or several items. Food stimuli appeared partially unwrapped with some of the product (either one or several pieces) appearing outside of the packaging. Stimuli appeared on the screen for a fixed duration of 3.6 seconds. Participants were asked to respond within that time limit and their responses were highlighted from the time they made a response until the end of the 3.6 second window, when the stimuli disappeared from the screen. Stimulus presentations were separated by a fixed inter-stimulus interval (ISI) of 6 seconds consisting of a central fixation cross. Each of the two scan runs consisted of 36 trials lasting five minutes and fifty seconds.

2.2.4. Training—The cue-approach training task was developed by Schonberg et al. (2014a). For each trial, images of the food items were presented on the screen for 1.2 seconds followed by a fixed ISI of 3.6 seconds (Figure 2B). Item order was randomized within a block of 32 trials. Participants were instructed to press a button on the keypad as fast as possible only when they heard an infrequent neutral tone and before the item disappeared from the screen. Items that were assigned to the Go condition were consistently associated with the tone. The tone appeared on average 950 ms after the item was presented on the screen (Go-signal delay, GSD). GSD was adjusted using a ladder technique. We increased the GSD by 17 ms if participants pressed the button before the item disappeared (to make the task more difficult) and reduced GSD by 50 ms if the participant failed to press the button or pressed it after the item disappeared (to make the task easier). We chose this 3:1 ladder titration ratio to ensure a 75% success rate in correct button presses. All 32 food items used during training were presented 12 times each during training. Each of the six scan runs consisted of two presentations of each stimulus (i.e. 64 trials) lasting five minutes and twelve seconds.

2.2.5. Probe—At the end of training, participants filled out a computer-adapted version of the Barratt impulsiveness scale questionnaire (BIS-11, Patton et al., 1995) while undergoing

a structural scan. They were then presented with pairs of food items in a probe task (Figure 2D). Items in each pair were matched for WTP and made up of one Go and one NoGo item (Figure 1B). Participants were told that a single trial would be drawn at random at the end of the session and their choice on that trial would be honored (i.e. they would receive the item that they chose on the randomly selected trial at the end of the experiment and remain in the lab to consume it). Pairs of items were presented on the screen for 2.4 seconds. Item selection was confirmed with a green rectangle drawn around the selected item (see Figure 2D), which remained on the screen from response time to the end of the 2.4 second trial window. If participants failed to make a choice within two seconds, a brief message asking them to respond faster appeared for 400 ms. Consecutive stimulus presentations were separated by a fixed ISI of 3.6 seconds. Each unique pair of items was presented in random order twice during probe (i.e. 64 probe trials total) for a scan run duration of six minutes and twenty-four seconds.

2.3. fMRI Acquisition

Imaging data were acquired on a 3 T Siemens Skyra MRI scanner with a 32-channel head coil. Functional data were acquired using a T2*-weighted multiband echo planar imaging sequence (repetition time (TR) = 1200 ms, echo time (TE) = 30 ms, flip angle (FA) = 63, field of view (FOV) = 230 mm, acquisition matrix of 96 × 96). Sixty four oblique axial slices were acquired with a 2.4 mm in-plane resolution positioned 30° off the anterior commissure-posterior commissure line to reduce the frontal signal dropout (Deichmann et al., 2003) and spaced 2 mm with a 0.4 mm gap to achieve full brain coverage. Slices were acquired using the multi-band sequence (Moeller et al., 2010) in an interleaved fashion. Each of the localizer runs consisted of 292 volumes, each of the training runs consisted of 260 volumes, and the probe run consisted of 324 volumes. In addition to functional data, a single three-dimensional high-resolution full brain image was acquired using a magnetization prepared rapid gradient echo (MPRAGE) pulse sequence (TR = 2400 ms, TI = 1000 ms, TE = 1.94 ms, FA = 8, FOV = 205 mm, voxel size = 0.8 × 0.8 × 0.8 mm) for brain masking and image registration.

2.4. Analysis

2.4.1. Behavioral Analysis

2.4.1.1. Probe: To test whether cue-approach training induced a preference change, we performed repeated-measures logistic regression to compare the odds of choosing the Go to NoGo items against equal odds for the high-value and low-value pairs separately. We also performed repeated-measures linear regression to test for differences in reaction time (RT) for choices of Go and NoGo items for the high-value and low-value pairs separately.

2.4.1.2. Auction: We ran repeated-measures linear regression to test the two-way interaction between time (pre-training/post-training auction) and condition (Go/NoGo) on WTP within high-value and low-value items separately. This interaction tests whether the change in WTP over time is different for Go and NoGo items. P values for the effects in the mixed models were calculated using the Kenward-Roger approximation for degrees of freedom (Kenward and Roger, 1997).

2.4.2. Imaging Analysis

2.4.2.1. Imaging Data Preprocessing: Raw imaging data in DICOM format were converted to NIFTI format and preprocessed through a standard preprocessing pipeline using the FSL package version 5 (Smith et al., 2004). Functional image time series were first aligned using the MCFLIRT tool to obtain six motion parameters that correspond to the x-y-z translation and rotation of the brain over time. Second, the skull was removed from the T2* images using the brain extraction tool (BET) and from the high-resolution T1 images using Freesurfer (Dale et al., 1999; Ségonne et al., 2004). Spatial smoothing was performed using a Gaussian kernel with a full-width half maximum (FWHM) of 5 mm. Data and design matrix were high-pass filtered using a Gaussian-weighted least-squares straight line fit with a cutoff period of 100 s. Grand-mean intensity normalization of each run's entire four-dimensional data set by a single multiplicative factor was also performed. The functional volumes for each participant and run were registered to the high resolution T1-weighted structural volume using a boundary-based registration method implemented in FSL5 (BBR, Greve and Fischl, 2009). The T1-weighted image was then registered to the MNI152 2 mm template using a linear registration implemented in FLIRT (12 degrees of freedom). These two registration steps were concatenated to obtain a functional-to-standard space registration matrix.

2.4.2.2. Cognitive Localizer: We conducted a GLM analysis on the cognitive localizer task data. The GLM model included eight regressors of interest: (i) onsets for valuation trials, modeled with a duration which equaled the average RT across all trials and participants; (ii) same onsets and duration as i but modulated by response (1 for like least to 4 for like most) demeaned across these trials within each run for each participant; (iii) onsets for perceptual decision trials modeled with the same duration as for i; (iv) same onsets and duration as iii but modulated by response (1 for single item and 2 for several items outside of packaging) demeaned across these trials within each run for each participant; (v) onsets for memory retrieval trials modeled with the same duration as for i; (vi) same onsets and duration as v but modulated by response (1 never saw this item in a store to 4 seen it within the last week) demeaned across these trials within each run for each participant; (vii) to account for any differences in RT between trial types we added a regressor with the onsets of all valid trials and the same duration as all other regressors (average RT across all trials and participants), while the modulator was the demeaned RT across all valid trials; (viii) onsets for missed trials. We included the six x, y, z translation and rotation motion parameters obtained from MCFLIRT, framewise displacement (FD) and RMS intensity difference from one volume to the next (DVARS, Power et al., 2012) as confound regressors. We also modeled out volumes with FD and DVARS that exceeded a threshold of 0.5 by adding a single time point regressor for each "to-be-scrubbed" volume (Siegel et al., 2013). All regressors were entered at the first level of analysis and all (but the added confound regressors) were convolved with a canonical double-gamma hemodynamic response function. The temporal derivative of each regressor (except for the added confound regressors) was included in the model. The model was estimated separately for each participant and each run. The raw parameter estimates from this model were used to train a classifier to decode the three localizer task conditions (i.e. perceptual decision, memory retrieval, and valuation).

2.4.2.3. Training: The GLM during the training phase included 4 regressors for each of Go and NoGo trial types broken down by the two subsequent probe trial types (high-value Go versus high-value NoGo and low-value Go versus low-value NoGo): (i) onsets of the Go trial, modeled with a fixed duration of 1.2 seconds; (ii) same onset and duration as i but modulated by subsequent number of times chosen during probe; (iii) same onset and duration as i but modulated by initial WTP; (iv) same onset and duration as i but modulated by the Go-signal delay for that trial. Thus there were two different Go trials (high and low) and for each there were four regressors yielding a total of 8 regressors. Then for each of the different types of NoGo trials there were three regressors similar to above except for modulation by Go signal delay as there was no go-signal in the NoGo trials. There were two different NoGo trial types (high- and low-value) and for each there were three regressors, thus yielding a total of 6 additional regressors. Additionally, for each high-value and low-value item that was not used during probe, we included the equivalent to regressors i and iii above to yield four additional regressors. To account for RT differences between all trials we added a regressor with the onsets of all Go trials and the modulator was the demeaned RT across all these trials. We further added a missed trial regressor each for high-value Go and low-value Go as well as two regressors for an erroneous response for high-value and low-value NoGo trials. There were a total of 23 regressors. We added the same covariates as in the probe design matrix, including the six motion regressors described above, along with FD and DVARS as confound regressors. This analysis did not yield significant results in our original study, but we found that we had increased power in the current study.

2.4.2.4. Multivariate Pattern Analysis: Localizer task: We performed a multivariate pattern analysis (MVPA) to classify the pattern of activation during each of the three cognitive processes engaged during the three trial types in the cognitive localizer task. We used whole brain raw parameter estimates obtained from the cognitive localizer task GLM described above as input into a three-class SVM classifier to classify the pattern for each of valuation, perceptual decision and memory retrieval cognitive processes. We then conducted two-way cross validation, where we trained the classifier on the first half of the cognitive localizer neural data and tested it on the second half (different runs), then vice versa to obtain average classifier cross-validation accuracy.

Cue-approach training: Once we ascertained that the cross-validation accuracy surpassed chance classification (i.e. within-localizer task cross-validation accuracy significantly above 33%), we trained the classifier on all the neural data and applied the classifier to raw parameter estimates extracted from the GLM on cue-approach training task data. We obtained classifier evidence scores (i.e. the classifier's estimate of the match between the test pattern and the trained patterns) for the valuation, perceptual-decision and memory-retrieval processes on each cue-approach training trial per participant. We ran this analysis in order to study changes in classifier evidence corresponding to each of the cognitive processes thought to be engaged during cue-approach training and to determine whether the increase in memory-retrieval and valuation classifier evidence across training predicts later choice at probe.

2.4.2.5. Probe: In line with the work we had previously carried out exploring the neural signature of value change following cue-approach training (Schonberg et al., 2014a), we focused our univariate analysis in the current study on the probe phase. We used a general linear model (GLM) for the probe phase that included seven regressors for each of the two trial types. For high-value Go versus high-value NoGo, (i) onsets of trials when high-value Go items were chosen with fixed duration, which was the average RT across all trials and participants; (ii) to explore the preference for each item, we used the demeaned total number of choices (on all probe trials where this item appeared) for the chosen item as a parametric modulator of the above onset regressor, with the same average RT as above used for duration; (iii) to account for the difference in pre-training WTP between the items in each pair we added the WTP difference as a parametric modulator with the same onsets and durations as regressor i. All of the above three regressors were added for trials when participants chose the NoGo item in a pair. To account for RT differences between choices of the Go and NoGo items we added a regressor with the onsets of all high-value Go and NoGo trials but as the modulator we added the demeaned RT across all these trials. We defined the same seven regressors for the probe trials that compared low-value Go to low-value NoGo, which resulted in a total of 15 regressors (two trial types times seven) and an additional regressor for missed trials of all types. The same motion, FD and DVARS confound regressors described above were included.

To test which regions showed greater modulation by preference for an item, we contrasted the parametric modulator of the chosen high-value Go items (regressor (ii) above) with the same regressor for the high-value NoGo items. We masked this contrast by our a priori anatomical mPFC region. The mask was the same as that previously used in Schonberg et al. (2014a) and encompassed the medial PFC by combining Harvard-Oxford regions (frontal pole, frontal medial cortex, paracingulate gyrus and subcallosal cortex) falling between $x = 14$, $y = -14$ and $z < 0$.

Ten participants were excluded from the imaging analysis because their parametric modulator of choices was a vector of zeroes. Two chose all high-value Go items in exactly the same proportions and three chose all high-value NoGo items in the same proportions during probe. One participant chose all low-value Go and four others chose all low-value NoGo items in exactly the same proportions. We mean-centered the choice regressors, resulting in a column of zeroes when they chose items in a particular category (e.g. high-value Go items) the same number of times. Thus, the parametric modulator was perfectly correlated with the intercept regressor (column of ones) resulting in a rank-deficient design matrix. For all group analyses we averaged across individual participants by performing a one-sample t-test to obtain the overall effects for the group. All reported statistical maps were corrected at the whole-brain level using a cluster-based Gaussian random field correction for multiple comparisons, with an uncorrected cluster-forming threshold of $z = 2.3$ and corrected extent threshold of $p < 0.05$, except for the comparison between preference modulation of Go and NoGo during probe, which was small volume corrected only for the anatomical mPFC mask (as was used in Schonberg et al., 2014a).

3. Results

3.1. Behavioral Results

Consistent with previous findings, we found an effect of cue-approach training on choices during the probe phase (Figure 3A). Participants chose high-value Go over high-value NoGo items on 65% of trials (odds ratio (O.R) = 2.21, 95% Confidence Interval (C.I) = [1.48 3.29], p < 0.0001). Also consistent with some previous findings when cue-approach training included a reduced stimulus set as is the case in this study, participants chose low-value Go over low-value NoGo items on 60% of trials (O.R = 1.7, C.I = [1.13 2.56], p = 0.01). However, the Go choice effect was larger for high-value than for low-value pairs (O.R = 1.28, C.I = [1.06 1.55], p = 0.01).

We repeated the initial auction after probe to test whether the subjective value placed on individual items changed after training. Although we had previously reported evidence that cue-approach training influenced the value of individual items, we did not replicate that finding in this study. WTP for high-value Go and NoGo items regressed equally toward the mean and WTP for low-value Go and NoGo items also increased equally and regressed toward the mean. There was a main effect of time (pre- to post-training, p < 0.0001), but no main effect of training condition (Go or NoGo) or interaction between the factors on WTP (p's > 0.5).

3.2. Imaging Results

3.2.1. MVPA Results—Localizer: Whole-brain pattern classifiers reliably distinguished fMRI activity (separately for each participant) from trials designed to elicit the cognitive processes of valuation, perceptual decision making, and memory retrieval of appetitive food items. We obtained above-chance classification accuracy (65.8% on average; chance is 33%) for each class. Thus, the classifiers can be applied to cue-approach task data to estimate the engagement of these three cognitive processes during training.

Cue-approach training: After training the classifier on all the cognitive localizer task fMRI data, we applied the classifier to fMRI data acquired during the cue-approach training task to predict the extent to which each of the cognitive processes of interest were elicited across training trials. In this analysis, we did not obtain any notable increases in the estimates across training (i.e. relatively flat lines in Figure 4 and no main effect of repetition number on classifier evidence for any of the three classes). In a mixed-effects linear regression model testing the interaction between repetition number (i.e. x-axis on plots in Figure 4) and Go status (Go solid green lines vs. NoGo dashed red lines in Figure 4) on valuation classifier evidence, we found no significant interaction and no main effects. Testing the same interaction on the memory classifier evidence, we found a trend-level interaction (β = −0.002, C.I = [−0.0054 0.0003], p = 0.08) and a main effect of Go status (β = −0.023, C.I = [−0.033 −0.013], p < 0.0001), but no main effect of repetition number on memory classifier evidence. Finally, testing the interaction on perceptual classifier evidence, we found a trend-level interaction (β = 0.003, C.I = [−0.0002 0.006], p = 0.06) and a main effect of go status (β = 0.015, C.I = [0.004 0.026], p = 0.006), but no effect of repetition number on perceptual classifier evidence.

Given that classification of cognitive processes did not vary meaningfully across the training period, we tested whether we at least had above-chance cross-task classification accuracy for the valuation process, which should be central to the cue-approach task. That is, we assessed whether valuation processes engaged during the localizer task were also engaged to some extent, and could be identified by the classifier, in the training task. To obtain a measure of cross-task classification accuracy, we trained a classifier on the localizer task data to distinguish between low-, medium-, and high-valued items from the auction. We trained a classifier to distinguish patterns of activity during the localizer task elicited by three levels of value during valuation trials (low/medium/high value), then tested the classifier on the cue-approach task data. The within-localizer task two-way cross-validation accuracy for this value-level classifier was 61.9% (chance is 33%). However, the cross-task classification accuracy was only 31.24%, which was no different than chance accuracy. Thus, we did not achieve proper cross-task classification accuracy using the item value-level classifier. This suggests that valuation was not elicited similarly during the localizer task and the cue-approach training task. Despite non-satisfactory cross-task validation, we found that there was a significant trial-by-trial variance in classifier evidence during the training phase, we sought to test whether trial-by-trial variance explained behavior during the later probe phase.

In a mixed-effects linear regression model, there was no main effect of valuation classifier evidence from the last presentation during the training phase on subsequent choice during probe, but there was an interaction between valuation classifier evidence from the last presentation during training and item type (high-value Go/NoGo) on subsequent choice during probe (Figure 5, p = 0.03). This suggests that the item-by-item relationship between valuation classifier evidence on the last presentation during the training phase and subsequent choices during probe was different for high-value Go and NoGo items. This interaction effect did not hold for the first training presentation (there was no three-way interaction between valuation classifier evidence by item type [high-value Go/NoGo] by presentation number [first/last]). The interaction also did not hold for low-value Go vs. NoGo, and did not hold when using memory retrieval or perceptual decision classifier evidence.

3.2.2. Univariate Results

3.2.2.1. Localizer: We tested the parametric modulation of preference level for foods measured through a four-alternative forced choice (least to most) to the question "How much would you like to eat this item?" We found that this measure of value was related to BOLD activity primarily in vmPFC (Figure 6 and Table 1), in line with numerous previous studies that have demonstrated a role for vmPFC in coding subjective value (for a meta-analysis, see Bartra et al., 2013).

3.2.2.2. Probe: We used the number of times an item was chosen at probe as a parametric modulator to test whether the vmPFC represents value change during probe in our task. In line with our previous published study (Schonberg et al., 2014a), we limited our analysis to a large anatomical area within mPFC. There were no whole-brain corrected or small-volume corrected (SVC) results for modulation of vmPFC BOLD by post-training preference for

high-value Go items. Additionally, the relationship between preference and BOLD in the vmPFC did not differ for choices of high-value Go and high-value NoGo items.

We ran the same analysis on the low value pair trials, since there was a behavioral effect in both low-value and high-value pair trials in this sample (unlike in Schonberg et al., 2014a). In this analysis, we found an amplified BOLD signal modulation by preference for choices of low-value Go over low-value NoGo items when restricting the analysis to an extensive anatomical mask of mPFC (SVC, Figure 7 and Table 2). There was no significant effect within the vmPFC for the modulation of BOLD by choices of low-value Go or NoGo items in the whole brain analysis.

3.2.2.3. Training: Consistent with previous findings, there were no differences in the Go stimulus onset driven activations for the last run of training compared to the first run. We also used the same parametric modulator as in the probe phase (i.e. the number of times a particular item was chosen) to test for preference change related signals during the last run of training. There was no effect within the vmPFC for the modulation of BOLD response by preference for any of high- or low-value Go or NoGo items. The lack of effect here fails to replicate previous findings. Additionally, the relationship between BOLD and preference for high-value and low-value Go vs. NoGo were no different, replicating previous lack of findings. However, the change in modulation by choice preferences (number of times a particular item is later chosen at probe) over time (run 6 minus run 1) is stronger for Go than for NoGo items in a number of regions that include left dlPFC and vlPFC (Figure 8 and Tables 3 & 4). This analysis includes both high- and low-value item training trials since we found a behavioral effect in both trial types. This finding is novel and was not present in the original study. The increased trial count may account for the positive finding in the current sample compared to the original study.

4. Discussion

Shifting preferences is key to behavioral change. Focus has recently turned to target automatic cognitive processes to influence choice behavior (Marteau et al., 2012). Previous work by our group has established the cue-approach training task as a viable paradigm to influence choice behavior without reverting to effortful self-control and external reinforcement (Schonberg et al., 2014a). The full underlying neural mechanism responsible for a shift in preferences following cue-approach training remains unknown. In the current study, we sought to investigate neural changes during cue-approach training that predict subsequent choices using multivariate pattern analysis as well as traditional univariate techniques.

We designed a cognitive localizer task that engages three distinct cognitive processes assumed to be implicated during cue-approach training to influence subsequent choices: valuation, memory and perceptual processing. Previous behavioral findings support the involvement of these processes during CAT training (Bakkour et al., 2016). We built a classifier to distinguish between whole-brain patterns of activity elicited by these three cognitive processes of interest (for a review of decoding techniques applied to fMRI, see Norman et al., 2006; Tong and Pratte, 2012). This classifier would in theory then have

allowed us to obtain a measure of the degree to which each of these cognitive processes were engaged during cue-approach training. Any changes in the obtained classifier evidence measures could then be regressed against choices at probe to determine the contribution of changes in the involvement of these cognitive processes during training in shifting preferences. Previous research supported this design. Indeed, several studies have demonstrated robust classification of abstract cognitive processes such as intentions (Haynes et al., 2007; Soon et al., 2008), attention (Rosenberg et al., 2015) and valuation (Gross et al., 2014) to name only a few. Further demonstrating the power of MVPA classifiers, the three-class classification discriminating the cognitive processes of memory retrieval, valuation and perceptual processing performed well above chance in a standard two-way cross-validation applied to the localizer task fMRI data. However, to our surprise, this classifier did not generalize well from the localizer task to the cue-approach training task. To better quantify the cross-task generalization of this classifier technique in our data, we trained another classifier to differentiate the patterns of brain activation elicited by three levels of value (low, medium, and high) during valuation trials of the localizer task. This value classifier performed well above chance in standard cross-validation applied to the localizer task fMRI data (for a simple guide on this method, see Mur et al., 2009), consistent with previous reports of value representation in the cortex (Krajbich et al., 2009; McNamee et al., 2013). However, this classifier did not perform above chance when applied to the cue-approach training task fMRI data. It is worth noting that we had a measure of willingness-to-pay for all foods presented during the training task and thus accurately labeled each food stimulus in terms of its level of value in the same way we labeled the level of value for localizer task stimuli. This labeling allowed us to accurately determine cross-task classification accuracy. The lack of cross-task value classification suggests that valuation, assumed to be elicited during cue-approach training, is not expressed in a similar enough manner in neural activation patterns as during the localizer task. One possibility for this discrepancy is that the localizer and cue-approach tasks are very different. The localizer task asks participants to process images of food items and answer two or four alternative forced choice questions, whereas during the cue-approach training task, participants are asked to simply view images of food unless they hear a tone that cues them to press a button. It appears that the patterns of activity elicited by potentially shared cognitive processes do not overlap enough to solve the classification problem posed. The lack of cross-task classification in this study stands in contrast to other studies that have demonstrated strong cross-task classification (Eger et al., 2009; Etzel et al., 2008; Lewis-Peacock and Postle, 2008; Meyer et al., 2010; Shinkareva et al., 2011). The lack of cross-task value-level classification success was surprising given previous work demonstrating cross-domain value-level classification (Gross et al., 2014). The lack of cross-task classification in the current study calls for future studies aimed at better defining the conditions under which cross-task classification is possible.

Despite the lack of expected average classification differences between training task conditions, we partially replicated previous behavioral results, and obtained novel imaging findings. First and foremost, the current study replicates the now well-established cue-approach behavioral effect. In the probe phase, participants chose items that were previously associated with a cued button press during the training phase over items that were not associated with a cue but that were matched for pre-experimental preference (Figure 2A).

Several other studies from our group (Bakkour et al., 2016) were conducted in order to narrow down the possible mechanisms engaged during cue-approach training to cause a preference shift. These studies suggest that a cued motor response that requires sustained attention prior to the cue is necessary to induce a change in choice behavior. Indeed, a cue alone without a motor response or an uncued motor response alone during the training phase are not sufficient to bias choices. Furthermore, eliminating the go-signal delay and sounding the cue to make a motor response concurrently with the onset of food stimuli without delay during the training phase eliminated the choice effect at probe. Finally, requiring participants to make choices using eye movements rather than manual button presses revealed a preference shift following standard cue-approach training involving a cued manual motor response with the cue sounding after the food stimulus appears. This last result provides evidence suggesting that the choice shift was not calculated within manual or ocular motor circuits but rather that the shift in preference is likely due to modulation of more general value coding regions in the brain such as vmPFC. Current findings further bolster the claim that cue-approach training modulates subjective value of individual items. Although cross-task classification for the level of value (low/medium/high value) from localizer task to cue-approach training was not significant, we leveraged the trial-by-trial variance in value classifier evidence from the three-class (value/memory/percept) classifier and found that the relationship between value classifier evidence on the last presentation and the number of times each food was later chosen at probe differed for high-value Go and NoGo items (Figure 5). Taking these results together, we suggest that cue-approach training engages attentional mechanisms during behaviorally relevant points in time in order to modulate value coding of items that were associated with the Go signal.

Previous imaging findings point to a more positive relationship between BOLD activation in the vmPFC and preference for choices of high-value Go when compared to choices of high-value NoGo items (Schonberg et al., 2014a). This finding was not replicated in the current study. Failure of replication could be due to low power. Our complete sample included 32 participants, which is considered adequate for fMRI studies. However, our contrasts of interest were based on participant behavior and we had to exclude participants who chose items within a category at the same rate, which represented a larger proportion of the sample than in a previous imaging study, significantly reducing our power to detect an effect in this analysis. Indeed, several participants chose items the same number of times during probe. This measure was entered as a parametrically modulated regressor in our probe phase GLM. In order to ensure that the parametric regressor is not correlated with the unmodulated regressor, we demeaned the choice measure entered into the parametric regressor, resulting in a column of zeroes when items were chosen the same number of times, and rendering the matrix rank deficient. We excluded ten participants from this analysis for this reason, reducing our sample size from 32 to 22 for the replication analysis. Furthermore, poor signal-to-noise ratio (SNR) in the vmPFC might be another reason for this lack of replication (Stenger, 2006). However, there was a strong relationship between ratings of how much participants wanted to eat an item during the cognitive localizer task and signal in the vmPFC (Figure 6), attesting to adequate SNR in that region. In our original imaging study, participants did not choose Go items more often that NoGo items in low-value pairs and we saw no difference in the modulation of the number of times lower-value items were chosen

during probe in the imaging analysis. However, in the current study, we found a behavioral effect in the low-value as well as the high-value pairs (Figure 3A). Consistent with the behavioral effect, the imaging analysis revealed that the modulation of probe phase BOLD by number of times a particular item was chosen was higher for low-value Go than for low-value NoGo items in the vmPFC (Figure 7). This suggests that the neural mechanism responsible for modulating item values during CAT is similar regardless of the initial value of the item.

Beyond the imaging findings during probe, we set out to look for neural changes during cue-approach training that might predict later choices, thus the motivation for the MVPA classifier analyses discussed above. Given that we did not achieve appropriate cross-task classification and that we found a behavioral effect in both low- and high-value pairs at probe, we pooled across both high- and low-value item trials during training for traditional univariate analyses on the training phase data. Consistent with previous findings, there were no differences in the Go stimulus onset-driven activations for the last run of training compared to the first run. This is likely due to the simple nature of the training phase and the fact that no decisions regarding the food were required during the training phase. The doubling of training trials for analysis (i.e. combining across high- and low-value item trials) increased our power to detect an effect (Liu and Frank, 2004) during training over our original study (Schonberg et al., 2014a). Changes in the relationship between BOLD and later preference from the beginning to the end of cue-approach training differed between Go and NoGo trials mainly in lateral prefrontal cortical areas (Figure 8). These areas are typically associated with task control (Dosenbach et al., 2007; 2006). Additionally, activity in the dlPFC has been shown to modulate the activity of the vmPFC (Hare et al., 2009), a region thought to code for subjective value (Bartra et al., 2013; Padoa-Schioppa and Assad, 2006). We suggest that task control increases for Go item trials as training progresses, given that Go food items are consistently associated with a tone and button press thus participants learn that they are in a Go trial when a Go item appears by the end of training. We also venture that task control does not change as training progresses for NoGo item trials. We further speculate that the increased task control subserved by dlPFC modulates subjective value of Go items coded in vmPFC.

Understanding the neural changes during non-reinforced training that underlie a later shift in choice preferences is important to the study of behavioral change. In the current study, we found that the extent to which a participant's whole-brain pattern of activity during a Go trial at the end of training reflected a valuation process determined the number of times he or she later chose that item. We also found that BOLD activity in a network of frontal regions was differentially related to later choices for Go and NoGo items as CAT training progressed. These findings could help us further increase the effectiveness of CAT training in nudging individuals to make lasting, positive changes in their choice behavior.

5. Conclusions

The cue-approach task continues to prove to be a useful paradigm for the study of behavioral change and potentially for the development of real-world interventions to help change and maintain habits. Its non-reliance on effortful self-control and its targeting of automatic

processes in the brain render it particularly appealing. This research shows promise for the development of new real-world, non-externally reinforced behavioral change paradigms by tapping attentional and memory mechanisms that act at behaviorally relevant points in time to modify valuation of particular stimuli. Training regimens that do not rely on external reinforcement, such as cue-approach training, could inspire the development of novel, lightweight behavioral treatments that help combat addiction, eating disorders and other ills by affecting lasting changes in choice behavior.

Acknowledgements

References

Bakkour A, Leuker C, Hover AM, Giles N, Poldrack RA, Schonberg T. Mechanisms of Choice Behavior Shift Using Cue-approach Training. Front Psychol. 2016; 7:421. doi:10.3389/fpsyg. 2016.00421. [PubMed: 27047435]

Bartra O, McGuire JT, Kable JW. The valuation system: a coordinate-based meta-analysis of BOLD fMRI experiments examining neural correlates of subjective value. Neuroimage. 2013; 76:412–427. doi:10.1016/j.neuroimage.2013.02.063. [PubMed: 23507394]

Becker GM, Degroot MH, Marschak J. Measuring utility by a single-response sequential method. Behav Sci. 1964; 9:226–232. doi:10.1002/bs.3830090304. [PubMed: 5888778]

Bjork, RA. Recency and recovery in human memory. In: Roediger, HL.Nairne, JS.Neath, I., Surprenant, AM., editors. The Nature of Remembering: Essays in Honor of Robert G. Crowder. American Psychological Association Press; Washington, DC: 2001. p. 211-232.

Bouton MEM. Context and behavioral processes in extinction. Learn Memory. 2004; 11:485–494. doi: 10.1101/lm.78804.

Bouton MEM. Context, time, and memory retrieval in the interference paradigms of Pavlovian learning. Psychol Bull. 1993; 114:80–99. [PubMed: 8346330]

Cahill K, Perera R. Competitions and incentives for smoking cessation. Cochrane Database Syst Rev. 2011:CD004307. doi:10.1002/14651858.CD004307.pub4. [PubMed: 21491388]

Cox DD, Savoy RL. Functional magnetic resonance imaging (fMRI) "brain reading": detecting and classifying distributed patterns of fMRI activity in human visual cortex. Neuroimage. 2003; 19:261–270. [PubMed: 12814577]

Dale AM, Fischl B, Sereno MI. Cortical surface-based analysis. I. Segmentation and surface reconstruction. Neuroimage. 1999; 9:179–194. doi:10.1006/nimg.1998.0395. [PubMed: 9931268]

De Martino B, Kumaran D, Seymour B, Dolan RJ. Frames, biases, and rational decision-making in the human brain. Science. 2006; 313:684–687. doi:10.1126/science.1128356. [PubMed: 16888142]

Deichmann R, Gottfried JA, Hutton C, Turner R. Optimized EPI for fMRI studies of the orbitofrontal cortex. Neuroimage. 2003; 19:430–441. [PubMed: 12814592]

Dosenbach NUF, Fair DA, Miezin FM, Cohen AL, Wenger KK, Dosenbach RAT, Fox MD, Snyder AZ, Vincent JL, Raichle ME, Schlaggar BL, Petersen SE. Distinct brain networks for adaptive and stable task control in humans. Proc Natl Acad Sci USA. 2007; 104:11073–11078. doi:10.1073/pnas.0704320104. [PubMed: 17576922]

Dosenbach NUF, Visscher KM, Palmer ED, Miezin FM, Wenger KK, Kang HC, Burgund ED, Grimes AL, Schlaggar BL, Petersen SE. A core system for the implementation of task sets. Neuron. 2006; 50:799–812. doi:10.1016/j.neuron.2006.04.031. [PubMed: 16731517]

Eger E, Michel V, Thirion B, Amadon A, Dehaene S, Kleinschmidt A. Deciphering cortical number coding from human brain activity patterns. Curr. Biol. 2009; 19:1608–1615. doi:10.1016/j.cub. 2009.08.047. [PubMed: 19781939]

Esterman M, Chiu Y-C, Tamber-Rosenau BJ, Yantis S. Decoding cognitive control in human parietal cortex. Proc Natl Acad Sci USA. 2009; 106:17974–17979. doi:10.1073/pnas.0903593106. [PubMed: 19805050]

Etzel JA, Gazzola V, Keysers C. Testing simulation theory with cross-modal multivariate classification of fMRI data. PLoS ONE. 2008; 3:e3690. doi:10.1371/journal.pone.0003690. [PubMed: 18997869]

Greve DN, Fischl B. Accurate and robust brain image alignment using boundary-based registration. Neuroimage. 2009; 48:63–72. doi:10.1016/j.neuroimage.2009.06.060. [PubMed: 19573611]

Gross J, Woelbert E, Zimmermann J, Okamoto-Barth S, Riedl A, Goebel R. Value signals in the prefrontal cortex predict individual preferences across reward categories. J Neurosci. 2014; 34:7580–7586. doi:10.1523/JNEUROSCI.5082-13.2014. [PubMed: 24872562]

Hare TA, Camerer CF, Rangel A. Self-control in decision-making involves modulation of the vmPFC valuation system. Science. 2009; 324:646–648. doi:10.1126/science.1168450. [PubMed: 19407204]

Haxby JV, Gobbini MI, Furey ML, Ishai A, Schouten JL, Pietrini P. Distributed and overlapping representations of faces and objects in ventral temporal cortex. Science. 2001; 293:2425–2430. doi:10.1126/science.1063736. [PubMed: 11577229]

Haynes J-D, Sakai K, Rees G, Gilbert S, Frith C, Passingham RE. Reading hidden intentions in the human brain. Curr. Biol. 2007; 17:323–328. doi:10.1016/j.cub.2006.11.072. [PubMed: 17291759]

Higgins ST, Budney AJ, Bickel WK, Badger G, Foerg FE, Ogden D. Outpatient behavioral treatment for cocaine dependence: One-year outcome. Exp Clin Psychopharmacol. 1995; 3:205–212.

Houben K, Havermans RC, Nederkoorn C. Beer à no-go: learning to stop responding to alcohol cues reduces alcohol intake via reduced affective associations rather than increased response inhibition - Houben - 2012 - Addiction - Wiley Online Library. Addiction. 2012

Kenward MG, Roger JH. Small sample inference for fixed effects from restricted maximum likelihood. Biometrics. 1997; 53:983–997. doi:10.2307/2533558. [PubMed: 9333350]

Krajbich I, Camerer C, Ledyard J, Rangel A. Using neural measures of economic value to solve the public goods free-rider problem. Science. 2009; 326:596–599. doi:10.1126/science.1177302. [PubMed: 19745115]

Lawrence NS, Verbruggen F, Morrison S, Adams RC. Stopping to food can reduce intake. Effects of stimulus-specificity and individual differences in dietary restraint. Appetite. 2015

Lemm S, Blankertz B, Dickhaus T, Müller K-R. Introduction to machine learning for brain imaging. Neuroimage. 2011; 56:387–399. doi:10.1016/j.neuroimage.2010.11.004. [PubMed: 21172442]

Lenartowicz A, Verbruggen F, Logan GD, Poldrack RA. Inhibition-related activation in the right inferior frontal gyrus in the absence of inhibitory cues. J Cognitive Neurosci. 2011; 23:3388–3399.

Lewis-Peacock JA, Drysdale AT, Oberauer K, Postle BR. Neural evidence for a distinction between short-term memory and the focus of attention. J Cognitive Neurosci. 2012; 24:61–79. doi:10.1162/jocn_a_00140.

Lewis-Peacock JA, Drysdale AT, Postle BR. Neural Evidence for the Flexible Control of Mental Representations. Cereb Cortex. 2015; 25:3303–3313. doi:10.1093/cercor/bhu130. [PubMed: 24935778]

Lewis-Peacock JA, Postle BR. Temporary activation of long-term memory supports working memory. J Neurosci. 2008; 28:8765–8771. doi:10.1523/JNEUROSCI.1953-08.2008. [PubMed: 18753378]

Lin JY, Pype AD, Murray SO, Boynton GM. Enhanced memory for scenes presented at behaviorally relevant points in time. PLoS Biol. 2010; 8:e1000337. doi:10.1371/journal.pbio.1000337.g004. [PubMed: 20305721]

Liu TT, Frank LR. Efficiency, power, and entropy in event-related FMRI with multiple trial types: Part I: theory. Neuroimage. 2004

Logan GD, Cowan WB. On the ability to inhibit thought and action: A theory of an act of control. Psychol Rev. 1984; 91:295–327. doi:10.1037/0033-295X.91.3.295.

Mahmoudi A, Takerkart S, Regragui F, Boussaoud D, Brovelli A. Multivoxel pattern analysis for FMRI data: a review. Comput Math Methods Med. 2012; 2012:961257. doi:10.1155/2012/961257. [PubMed: 23401720]

Marteau TM, Hollands GJ, Fletcher PC. Changing human behavior to prevent disease: the importance of targeting automatic processes. Science. 2012; 337:1492–1495. doi:10.1126/science.1226918. [PubMed: 22997327]

McNamee D, Rangel A, O'Doherty JP. Category-dependent and category-independent goal-value codes in human ventromedial prefrontal cortex. Nat Neurosci. 2013; 16:479–485. doi:10.1038/nn. 3337. [PubMed: 23416449]

Meyer K, Kaplan JT, Essex R, Webber C, Damasio H, Damasio A. Predicting visual stimuli on the basis of activity in auditory cortices. Nat Neurosci. 2010; 13:667–668. doi:10.1038/nn.2533. [PubMed: 20436482]

Moeller S, Yacoub E, Olman CA, Auerbach E, Strupp J, Harel N, Ugurbil K. Multiband multislice GE-EPI at 7 tesla, with 16-fold acceleration using partial parallel imaging with application to high spatial and temporal whole-brain fMRI. Magnet Reson Med. 2010; 63:1144–1153. doi:10.1002/mrm.22361.

Mur M, Bandettini PA, Kriegeskorte N. Revealing representational content with pattern-information fMRI--an introductory guide. Soc Cogn Affect Neurosci. 2009; 4:101–109. doi:10.1093/scan/nsn044. [PubMed: 19151374]

Norman KA, Polyn SM, Detre GJ, Haxby JV. Beyond mind-reading: multi-voxel pattern analysis of fMRI data. Trends Cogn Sci. 2006; 10:424–430. doi:10.1016/j.tics.2006.07.005. [PubMed: 16899397]

O'Doherty J, Dayan P, Schultz J, Deichmann R, Friston K, Dolan RJ. Dissociable roles of ventral and dorsal striatum in instrumental conditioning. Science. 2004; 304:452–454. doi:10.1126/science. 1094285. [PubMed: 15087550]

Padoa-Schioppa C, Assad JA. Neurons in the orbitofrontal cortex encode economic value. Nature. 2006; 441:223–226. doi:10.1038/nature04676. [PubMed: 16633341]

Patton JH, Stanford MS, Barratt ES. Factor structure of the Barratt impulsiveness scale. J Clin Psychol. 1995; 51:768–774. [PubMed: 8778124]

Plassmann H, O'Doherty J, Rangel A. Orbitofrontal cortex encodes willingness to pay in everyday economic transactions. J Neurosci. 2007; 27:9984–9988. doi:10.1523/JNEUROSCI.2131-07.2007. [PubMed: 17855612]

Polyn SM, Natu VS, Cohen JD, Norman KA. Category-specific cortical activity precedes retrieval during memory search. Science. 2005; 310:1963–1966. doi:10.1126/science.1117645. [PubMed: 16373577]

Power JD, Barnes KA, Snyder AZ, Schlaggar BL, Petersen SE. Spurious but systematic correlations in functional connectivity MRI networks arise from subject motion. Neuroimage. 2012; 59:2142–2154. doi:10.1016/j.neuroimage.2011.10.018. [PubMed: 22019881]

Rosenberg MD, Finn ES, Constable RT, Chun MM. Predicting moment-to-moment attentional state. Neuroimage. 2015; 114:249–256. doi:10.1016/j.neuroimage.2015.03.032. [PubMed: 25800207]

Schonberg T, Bakkour A, Hover AM, Mumford JA, Nagar L, Perez J, Poldrack RA. Changing value through cued approach: an automatic mechanism of behavior change. Nat Neurosci. 2014a; 17:625–630. doi:10.1038/nn.3673. [PubMed: 24609465]

Schonberg T, Bakkour A, Hover AM, Mumford JA, Poldrack RA. Influencing food choices by training: evidence for modulation of frontoparietal control signals. J Cognitive Neurosci. 2014b; 26:247–268. doi:10.1162/jocn_a_00495.

Ségonne F, Dale AM, Busa E, Glessner M, Salat D, Hahn HK, Fischl B. A hybrid approach to the skull stripping problem in MRI. Neuroimage. 2004; 22:1060–1075. doi:10.1016/j.neuroimage. 2004.03.032. [PubMed: 15219578]

Shinkareva SV, Malave VL, Mason RA, Mitchell TM, Just MA. Commonality of neural representations of words and pictures. Neuroimage. 2011; 54:2418–2425. doi:10.1016/j.neuroimage.2010.10.042. [PubMed: 20974270]

Siegel JS, Power JD, Dubis JW, Vogel AC, Church JA, Schlaggar BL, Petersen SE. Statistical improvements in functional magnetic resonance imaging analyses produced by censoring high-motion data points. Hum Brain Mapp. 2013 doi:10.1002/hbm.22307.

Slovic P. The construction of preference. Am Psychol. 1995; 50:364.

Smith SM, Jenkinson M, Woolrich MW, Beckmann CF, Behrens TEJ, Johansen-Berg H, Bannister PR, De Luca M, Drobnjak I, Flitney DE, Niazy RK, Saunders J, Vickers J, Zhang Y, De Stefano N, Brady JM, Matthews PM. Advances in functional and structural MR image analysis and implementation as FSL. Neuroimage. 2004; 23(Suppl 1):S208–19. doi:10.1016/j.neuroimage. 2004.07.051. [PubMed: 15501092]

Soon CS, Brass M, Heinze H-J, Haynes J-D. Unconscious determinants of free decisions in the human brain. Nat Neurosci. 2008; 11:543–545. doi:10.1038/nn.2112. [PubMed: 18408715]

Stenger, VA. Technical considerations for BOLD fMRI of the orbitofrontal cortex. In: Zald, D., Rauch, S., editors. The Orbitofrontal Cortex. Oxford University Press; Oxford: 2006. doi:10.1093/acprof:oso/9780198565741.001.0001

Swallow KM, Jiang YV. The attentional boost effect: transient increases in attention to one task enhance performance in a second task. Cognition. 2010; 115:118–132. doi:10.1016/j.cognition. 2009.12.003. [PubMed: 20080232]

Thorndike, EL. Animal Intelligence: Experimental Studies. Macmillan; New York: 1911.

Tong F, Pratte MS. Decoding patterns of human brain activity. Annu Rev Psychol. 2012; 63:483–509. doi:10.1146/annurev-psych-120710-100412. [PubMed: 21943172]

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. doi:10.1111/j.1460-9568.2009.06796.x. [PubMed: 19490086]

Tversky A, Kahneman D. Rational choice and the framing of decisions. J Bus. 1986:S251–S278.

Veling H, Aarts H, Stroebe W. Stop signals decrease choices for palatable foods through decreased food evaluation. Front Psychol. 2013; 4:875. doi:10.3389/fpsyg.2013.00875. [PubMed: 24324451]

Verbruggen F, Logan GD. Automatic and controlled response inhibition: associative learning in the go/no-go and stop-signal paradigms. J Exp Psychol Gen. 2008; 137:649–672. doi:10.1037/a0013170. [PubMed: 18999358]

Wessel JR, O'Doherty JP, Berkebile MM, Linderman D, Aron AR. Stimulus devaluation induced by stopping action. J Exp Psychol Gen. 2014; 143:2316–2329. doi:10.1037/xge0000022. [PubMed: 25313953]

Wood W, Neal DT. A new look at habits and the habit-goal interface. Psychol Rev. 2007; 114:843–863. doi:10.1037/0033-295X.114.4.843. [PubMed: 17907866]

Yarkoni T, Poldrack RA, Nichols TE, Van Essen DC, Wager TD. Large-scale automated synthesis of human functional neuroimaging data. Nat Methods. 2011; 8:665–670. doi:10.1038/nmeth.1635. [PubMed: 21706013]
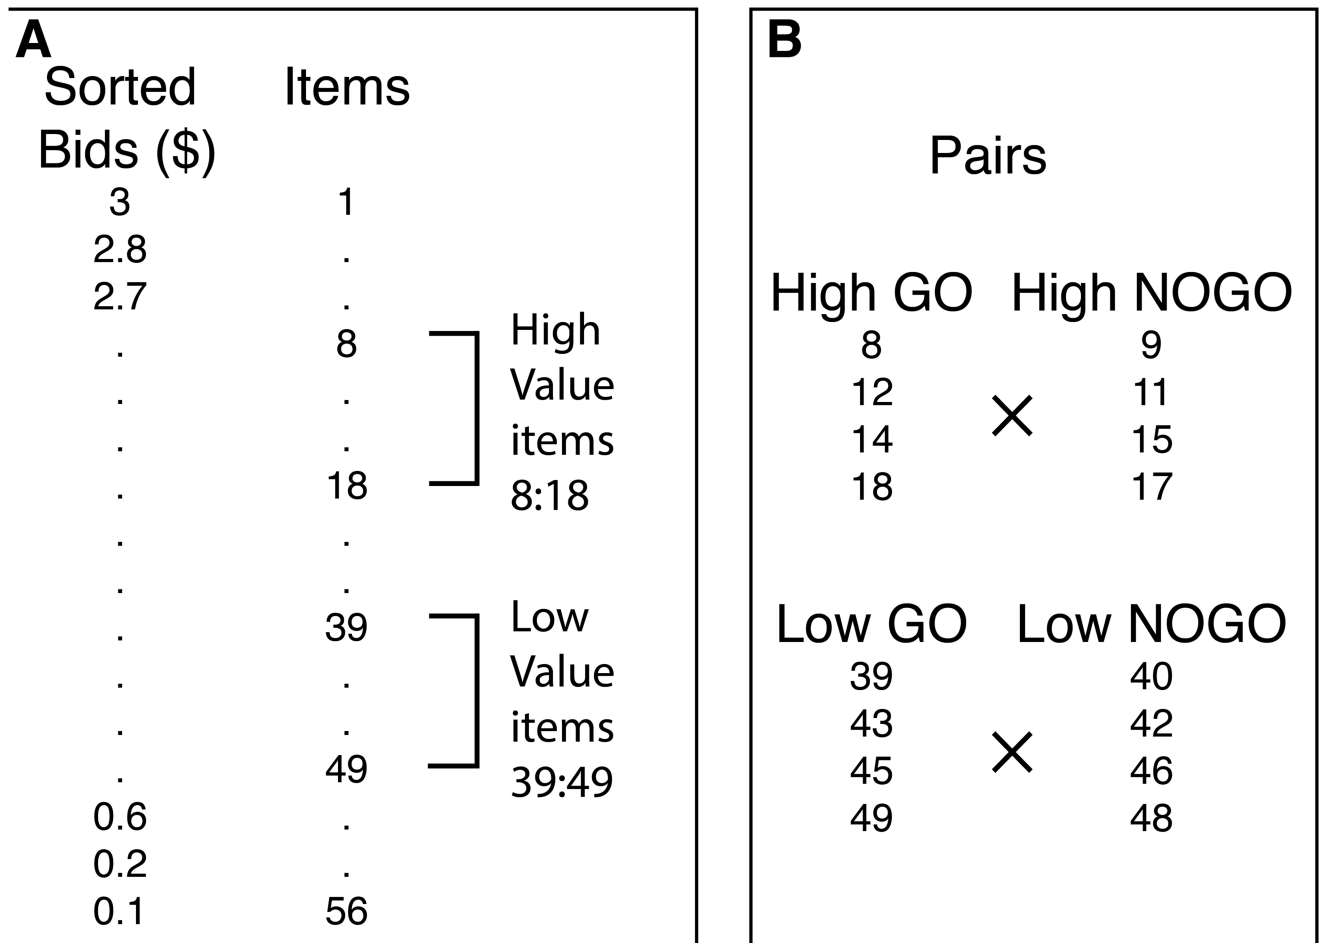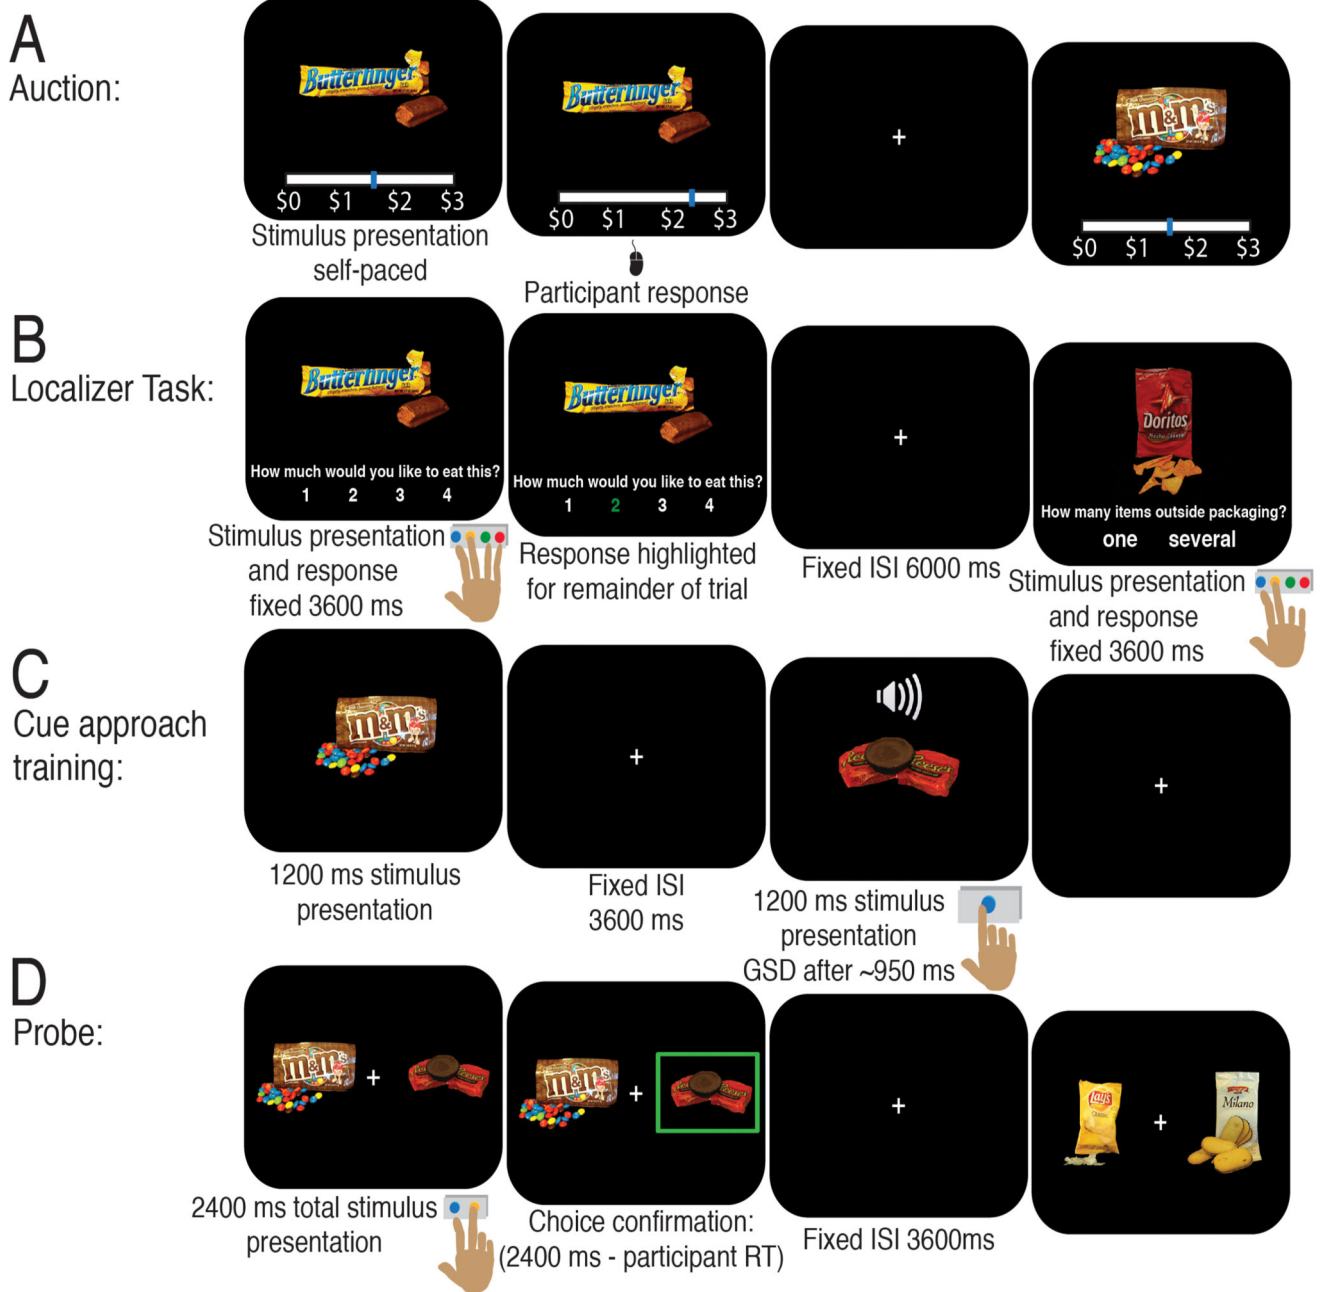
Figure 1.
Sorting and pair matching procedure. A) Items were rank ordered based on bid obtained in the auction (Figure 2A). Items were classified into high value (8:18) and low value items (39:49). B) 8 High and 8 low value items were assigned to one of two training conditions (Go, associated with a go-signal auditory cue and NoGo, not associated with a go-signal). Item Go/NoGo condition assignments were counterbalanced across participants.

Figure 2.
Task procedure. A) Auction. Participants placed their bid by selecting an amount between $0 and $3 using the mouse. B) Cognitive localizer task. Food items appeared on the screen one at a time, at the bottom of the screen one of three questions appeared. Participants were asked to answer the question relevant to the item on the screen within 3.6 seconds, at which time the trial ended. Successive stimulus presentations were separated by a fixed ISI of 6 s. C) Cue-approach training. Single food items appeared on the screen for a fixed 1.2 s. Participants were asked to press a button on the keypad as quickly as possible only when they heard a neutral tone that sounded on average 950 ms after food stimulus onset (GSD).

Stimulus presentations were separated by a fixed 3.6 s ISI. D) Probe task. Participants chose between two items on the screen. They were told that their choice on a random probe trial would be honored at the end of the experiment. Choices had to be made within 2 s of trial onset.
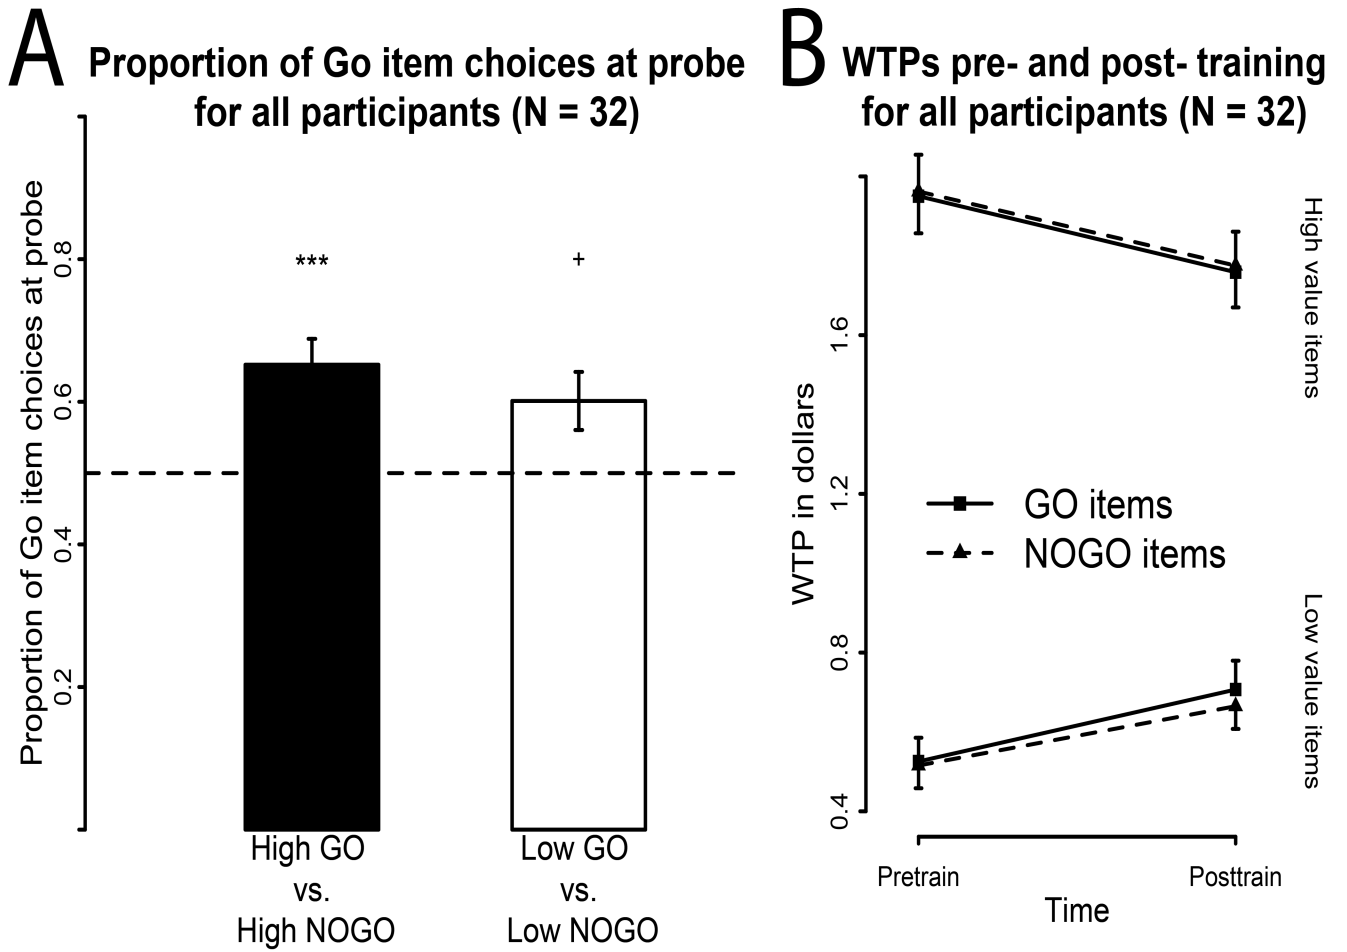
Figure 3.
Behavioral results for cue-approach study. A) Proportion of choices of the Go item in pairs of high-value Go versus NoGo and low-value Go versus NoGo items for all participants. Significance level reflects odds of choosing the Go to NoGo item. B) WTP before and after cue-approach training for Go and NoGo separately for items in the probe high-value Go versus high-value NoGo pairs (top) and low-value Go versus low-value NoGo pairs (bottom). The sample includes all participants. Error bars represent one standard error of the mean (SEM) in A and within-subject SEM in B. ***: $p < 0.0001$, +: $p < 0.05$ (two-sided tests).

Figure 4.
Average MVPA classifier evidence during cue-approach training. Three-class SVM classifier was trained on cognitive localizer neural data and applied to each trial of cue-approach training to obtain classifier evidence for each class split by training trial type (Go [solid green line] or NoGo [dashed red line]) for valuation (left), memory retrieval (middle) and perceptual decision (right). The shaded areas represent one within-participant SEM.
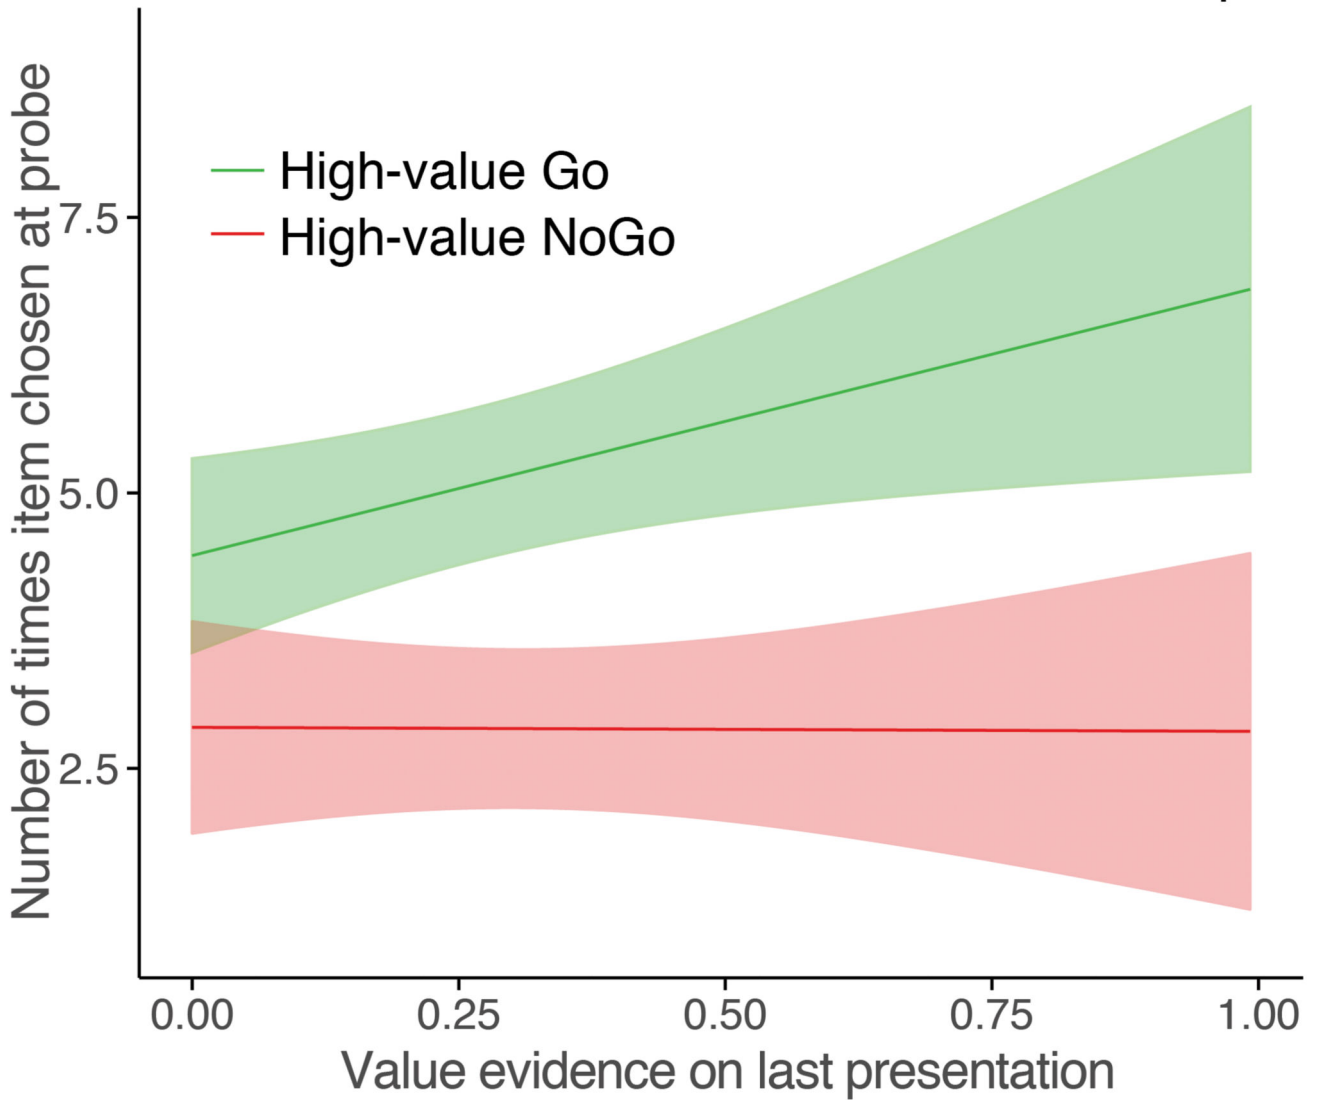
Figure 5.
Linear mixed-effects regression model interaction of item training type (high-value Go or NoGo) and value classifier evidence from the last cue-approach training trial on the number of times a particular item was later chosen at probe. The lines depict the group-level linear effects and the shaded areas depict 95% confidence curves.

Value level modulation during localizer task valuation trials

x = 4

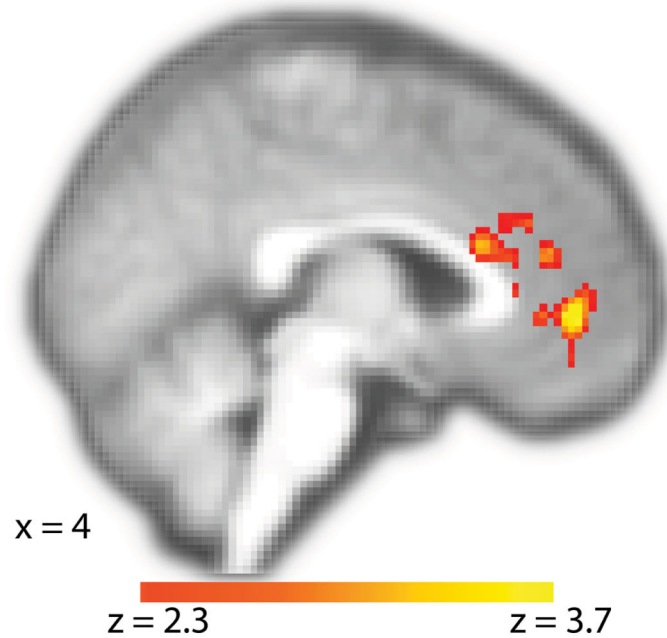z = 2.3   z = 3.7

Figure 6.
Localizer task imaging results (N=32) for the modulation by preference level measured by a four-alternative forced choice (least to most) to the question "How much would you like to eat this food?". Coordinates reported in standard MNI space. Heatmap color bars range from z-stat = 2.3 to 3.7. This map was cluster-corrected at a whole-brain level $p < 0.05$, two sided linear regression. To see the full map go to http://neurovault.org/images/24109/

Figure 7.
Probe imaging results (N=22) for the modulation by number of times a particular lower-value Go item was chosen greater than the modulation by number of times a particular lower-value NoGo item was chosen. Coordinates reported in standard MNI space. Heatmap color bars range from z-stat = 2.3 to 3.2. This map was cluster-corrected within an *a priori* defined anatomical mPFC mask (SVC) p < 0.05, two sided linear regression. To see the full SVC map go to http://neurovault.org/images/24110/

Figure 8.
Training imaging results (N=27). Changes in BOLD modulation by number of times a particular GO item was later chosen at probe from the beginning to the end of training grater than changes in BOLD modulation by number of times a particular NOGO item was later chosen at probe from beginning to the end of training. Coordinates reported in standard MNI space. Heatmap color bars range from z-stat = 2.3 to 3.2. This map was cluster-corrected at a whole-brain level p < 0.05, two sided linear regression. To see the full map go to http://neurovault.org/images/24111/

Table 1

Regions showing significant activations for the imaging contrast presented in Figure 6. The list shows all regions from the Harvard-Oxford atlas that contained more than 10 active voxels within the clusters, along with the peak x/y/z location for the cluster in MNI space.

Cluster Number	Region	# voxels in region	Cluster size	x	y	z	peak Z
1	L Lingual Gyrus	872					
	L Intracalcarine Cortex	431					
	L Occipital Fusiform Gyrus	178					
	L Occipital Pole	102	1989	−10	−74	−2	6.92
	R Lingual Gyrus	45					
	L Precuneous Cortex	12					
	L Supracalcarine Cortex	10					
2	L Paracingulate Gyrus	175					
	R Paracingulate Gyrus	116					
	L Frontal Pole	79	538	0	50	2	4.35
	L Cingulate Gyrus, anterior division	74					
	R Cingulate Gyrus, anterior division	42					
3	R Cingulate Gyrus, anterior division	171					
	L Cingulate Gyrus, anterior division	113	435	−2	36	20	3.98
	L Paracingulate Gyrus	80					
	R Paracingulate Gyrus	38					
4	R Angular Gyrus	156					
	R Middle Temporal Gyrus, temporooccipital part	155					
	R Supramarginal Gyrus, posterior division	73	430	64	−50	18	3.39
	R Lateral Occipital Cortex, superior division	16					
	R Lateral Occipital Cortex, inferior division	12					
5	L Supramarginal Gyrus, posterior division	241					
	L Angular Gyrus	65	382	−54	−44	48	4.05
	L Supramarginal Gyrus, anterior division	49					
	L Lateral Occipital Cortex, superior division	18					
6	R Inferior Frontal Gyrus, pars triangularis	161					
	R Inferior Frontal Gyrus, pars opercularis	50					
	R Insular Cortex	44					
	R Frontal Operculum Cortex	31	380	50	28	0	3.93
	R Frontal Pole	25					
	R Central Opercular Cortex	21					
	R Frontal Orbital Cortex	10					
7	L Middle Frontal Gyrus	244	378	−34	30	46	4.11
	L Superior Frontal Gyrus	80					

Table 2

Regions showing significant activations for the imaging contrast presented in Figure 7. The list shows all regions from the Harvard-Oxford atlas that contained more than 10 active voxels within the cluster, along with the peak x/y/z location for the cluster in MNI space.

Cluster Number	Region	# voxels in region	Cluster size	x	y	z	peak Z
1	L Frontal Pole	143	237	4	56	−6	3.7
	L Inferior Frontal Gyrus, pars triangularis	75					

Table 3

Regions showing significant activations for the imaging contrast presented in Figure 8. The list shows all regions from the Harvard-Oxford atlas that contained more than 10 active voxels within the clusters, along with the peak x/y/z location for the cluster in MNI space.

Cluster Number	Region	# voxels in region	Cluster size	x	y	z	peak Z
1	R Postcentral Gyrus	565	1365	10	−52	68	3.92
	R Superior Parietal Lobule	215					
	R Precentral Gyrus	193					
	R Precuneous Cortex	118					
	L Precentral Gyrus	85					
	R Supramarginal Gyrus, anterior division	47					
2	L Middle Frontal Gyrus	319	373	−42	12	42	3.46
	L Precentral Gyrus	26					
3	L Frontal Pole	143	258	−44	36	10	3.84
	L Inferior Frontal Gyrus, pars triangularis	75					

Table 4

Pearson correlations between the map in Figure 8 and formal term-based reverse inference maps using Neurosynth (Yarkoni et al., 2011)

Term	Correlation (r)
prefrontal	0.141
semantic	0.125
medial prefrontal	0.12
theory mind	0.11
default	0.107
mind	0.106
default mode	0.098
parietal	0.096
medial	0.094
mode	0.094