THE EFFECTS OF VALENCED FEEDBACK IN REAL-TIME fMRI REGULATION STUDIES

by

HEATHER ELIZABETH MANN

B.A., University of Waterloo, 2008

A THESIS SUBMITTED IN PARTIAL FULFILLMENT OF
THE REQUIREMENTS FOR THE DEGREE OF

MASTER OF ARTS

in

THE FACULTY OF GRADUATE STUDIES

(Psychology)

THE UNIVERSITY OF BRITISH COLUMBIA

(Vancouver)

August, 2010

Abstract

In real-time fMRI regulation studies, subjects view feedback showing the fluctuating activation within a particular region of their brains as they attempt to regulate that region's activation. This technology is of theoretical and clinical interest; however, it is unclear whether real-time regulation training is equally effective for all brain regions. Real-time feedback can be positive (if activation is in the desired direction) or negative (if in the opposite direction), suggesting a potential confound for training studies. We reasoned that if particular brain regions are differentially activated according to feedback valence, activations related to feedback might interact with the regulation task. Thus, we designed a study that allowed us to manipulate the valence of feedback in a real-time training context. Subjects were instructed to up-regulate and down-regulate their parahippocampal place area (PPA) in 30-second blocks while in the scanner, viewing feedback which they believed to reflect the activation of this region. In reality, the feedback was pre-constructed, and alternated between positive and negative valence blocks of varying length. Comparing positive with negative feedback, positive feedback activated nucleus accumbens, a reward centre, and certain emotion-relevant regions. Negative feedback produced little consistent activation over positive feedback. In general, feedback effects were greater for moderate feedback than strong feedback, possibly reflecting heightened uncertainty toward moderate feedback. We conclude that feedback-based activations are unlikely to interfere with regulation training for most cortical regions, though emotion-relevant regions may be more sensitive to feedback valence. We also propose that researchers explore feedback methods which emphasize reward-based learning.

Table of Contents

List of Figures

List of Tables

Acknowledgements

Thanks to my thesis advisor, Dr. Kalina Christoff, for her expert guidance and collaboration on this project. I would also like to thank my thesis committee, Dr. Alan Kingstone and Dr. Todd Handy, for their encouragement, support and thoughtful feedback.

My labmates have been wonderful to work with and to get to know over the course of this degree. I am grateful to Graeme McCaig for his patient guidance and his software development. Special thanks to Adi Steif for her proficient assistance with data analysis and inestimable motivation. Thanks to fellow grad students Melissa Ellamil and Matt Dixon for riding out the ups and downs together.

Friends and family were of much support to me during the course of this degree. I am grateful to all who have inspired me through their words and actions.

1 INTRODUCTION

For decades, training people to regulate their own physiological functions has been of interest to researchers and clinicians alike. The 1960s and '70s saw the development of electroencephalogram (EEG) neurofeedback, which measures synchronized electrical activity across the brain, and is now recognized as an empirically supported treatment for a wide variety of disorders (Yucha & Gilbert, 2004; Moss & Kirk, 2004). With the development of functional magnetic resonance imaging (fMRI) in the 1990s, researchers were able to non-invasively observe activations within specific anatomical brain regions, capitalizing on changing hemodynamics. Following the advent of this technology, researchers began to explore the possibility of creating real-time feedback from fMRI data as a means of training subjects to regulate their own localized brain activations.

Several research groups targeting various brain regions have now shown that presenting people with real-time fMRI feedback can help them improve at regulating target regions. However, while numerous regions have now been targeted for real-time training, much of the cortex remains uncharted territory. Furthermore, given the general bias against publishing null findings, real-time fMRI studies that do not yield significant training effects risk remaining in file drawers, never reaching the wider research community. One of the big unknowns of real-time fMRI regulation training is whether it can be successfully applied to all brain regions, or whether only certain regions are amenable to training (deCharms, 2007). The present work seeks to shed light on this query by examining the mechanism through which real-time training operates: fluctuating

visual feedback. Specifically, we wondered whether viewing feedback based on one's own regional brain activity might itself activate particular brain regions. If so, the presentation of real-time feedback would likely interact with regulation training for those regions.

1.1 Real-time fMRI Regulation Training and Its Applications

People regulate their brain activations all the time, via thoughts and actions (deCharms, 2007). However, individuals do not generally know or consider how activity within their brain coincides with particular mental processes or experiences. If we assume that the brain activity governs cognitive activity, a foundational tenet of cognitive neuroscience, then learning to systematically regulate brain activity should afford people additional control over their own mental processes. This represents the core principle of regulation training using real-time fMRI feedback.

Real-time regulation training involves presenting individuals with live information (feedback) derived from the blood oxygenation-dependent (BOLD) signal from a localized brain region. Often, people are given a cognitive strategy as a starting point for up-regulating and/or down-regulating the BOLD signal. Using real-time feedback (which usually carries a 1-3 second delay), subjects attempt to gain control over the BOLD activation signal by moving it in the desired direction.

1.1.1 Clinical applications of real-time regulation training

One of the primary reasons real-time regulation training has captured the attention of researchers is its potential applicability for treating clinical conditions. Most

suggestions for therapeutic applications are speculative, waiting to be bolstered by empirical evidence, but several conditions are of particular theoretical relevance. For instance, stroke often leads to loss of functioning in a particular brain region; real-time fMRI regulation training may provide stroke patients with an opportunity to regain functioning of the affected region (Weiskopf et al., 2007; Sitaram et al., 2007; deCharms et al., 2008). Psychopathy has been associated with underactivity of frontolimbic circuits (Kiehl, 2006); one research group has reported developing a real-time fMRI setup specifically designed for treating criminal psychopathy (Sitaram et al., 2007). In addition, several authors have suggested real-time regulation training as a potential treatment for emotion disorders (Weiskopf, Scharnowski et al., 2004; Sitaram et al., 2007; Johnston, Boehm, Healy, Goebel, & Linden, 2010).

To date, one study has offered convincing evidence that real-time fMRI feedback could function as a clinical treatment tool. deCharms and colleagues (2005) examined whether people could learn to control pain experiences using real-time feedback derived from the rostral anterior cingulate cortex (rACC), a region thought to be involved in pain perception. Participants were subjected to a painful thermal stimulus every 30 seconds, and instructed to intensify or diminish their pain experience in alternating blocks. The experimental group was shown feedback based on the BOLD signal in their rACC, while control groups were shown feedback from another brain region, feedback from a previous participant's rACC activations, no feedback, or were given twice as long to implement cognitive strategies. Not only did the experimental group show greater regulation of rACC with training than control groups, but the degree of change significantly correlated with changes in participants' ratings of the intensity and unpleasantness of their pain

experience. Impressively, a group of chronic pain patients who underwent a similar real-time training procedure with feedback from their rACC (but no external painful stimulus) reported increased control of their pain following training. For the chronic pain patients too, changes in pain ratings were correlated with changes in rACC regulation.

Given the promise of real-time training for treating clinical conditions, examining the mechanism behind it – the feedback itself – is worthwhile in arming us with better understanding of how this technology operates. I now turn to a wider review of published real-time training results.

1.2 Efficacy of Real-time Regulation Training

The first study to employ a continuously updating real-time fMRI feedback display (with a delay of less than two seconds) investigated a single subject's ability to regulate activation of his anterior cingulate cortex (ACC; Weiskopf et al., 2003). Every 60 seconds, this subject alternated between attempting to increase the ACC signal and letting it return to baseline; importantly, the signal change in rostral-ventral ACC increased over the course of training. Soon after, it was found that subjects who saw real-time feedback, but not controls, gained enhanced control over activation of the somatomotor cortex (deCharms et al., 2004). Another study with feedback indicating the differential activation of supplementary motor area (SMA) and parahippocampal place area (PPA) found training improvements in two of four subjects, suggesting individual variability in training responsiveness (Weiskopf, Mathiak et al., 2004).

More recent studies have investigated regulation training of emotion regions. Caria and colleagues (2007) used real-time feedback to train subjects to increase

activation in the right anterior insula. Nine subjects viewed live feedback while attempting to up-regulate this region by recalling personal, affective events; over three 4-minute feedback sessions, subjects achieved greater activation increases in their right anterior insula during up-regulation periods relative to rest. This improvement lasted in a post-training no feedback session, though the activation gain was not significant. Johnston and colleagues (2010) examined training effects for functionally-defined emotion regions – i.e. regions that were responsive to viewing negative emotional pictures, which included the amygdala, ventrolateral prefrontal cortex, and insula. Comparing activations during up-regulation periods with rest, they found that eleven of 13 subjects improved at activating their emotion regions. Notably, neither of these studies demonstrated down-regulation of emotion regions.

Another recent study focused on linguistic processing, targeting the right inferior frontal gyrus (IFG; Rota et al., 2009). Seven subjects viewed a feedback thermometer indicating activation of the right IFG during scanning, and alternated between blocks of up-regulation and rest. Over four training runs, these subjects learned to increase IFG, and showed subsequent improvement at an emotional prosody identification task.

Our lab investigated regulation training for two frontal brain regions: rostrolateral prefrontal cortex (rLPFC) and rostromedial prefrontal cortex (rMPFC). The investigation of rLPFC found that with feedback, subjects learned to alternate between up-regulation and down-regulation more effectively than control subjects who viewed another subject's feedback or received no feedback (Keramatian, 2009; McCaig, Dixon, Keramatian, Liu, & Christoff, in preparation). In spite of this group level trend, there was substantial

individual variability in training improvement, with some subjects' performance actually declining over the course of several feedback sessions.

Following interest in using real-time fMRI as a treatment tool for emotion disorders, our lab targeted rMPFC for real-time regulation training due to its putative role in emotional awareness (Lane, Fink, Chau, & Dolan, 1997; Ochsner et al., 2004). Subjects were instructed to regulate the rMPFC activation signal up and down, by reflecting on an emotional memory (for up-regulation) and focusing on the associated bodily sensations (for down-regulation). Across four scanning sessions with real-time feedback, four out of five subjects grew *worse* at modulating rMPFC. In addition, subjects produced greater regulation differences on a no feedback pre-training session than on a post-training session. The author speculated that "feedback methods which have previously been shown to be useful for helping subjects to learn to self-regulate brain activation may be limited in the extent to which they apply to other brain regions" (Smith, 2008).

Overall, while numerous studies have demonstrated training improvements with real-time feedback, at least one real-time fMRI investigation found that feedback had a detrimental effect on regulation performance.

1.3 The Role of Feedback in Real-Time Regulation Training

From the studies reviewed above, it appears that some regions are more amenable to real-time training than others. In addition, real-time regulation training has yet to be explored on a variety of cortical and subcortical brain areas.

In real-time regulation training studies, participants view feedback about their own performance, which fluctuates in strength and valence (positive or negative). If particular brain regions are responsive to positive or negative real-time feedback, regulation training for those regions may be compromised. To expand, if feedback valence modulates activation of a given region, showing real-time feedback as a part of regulation training may produce activations that interact with the regulation task. For instance, if subjects were instructed up-regulate a brain region activated by positive feedback, initial success at should facilitate further activation increases. However, if the task was to down-regulate (deactivate) the region, positive feedback should cause activation that makes the task more difficult. This effect may lead to an overall negative feedback cycle. Thus, valenced real-time feedback could conceivably interact with the regulation task in a helpful or harmful manner.

Another consideration is that if subjects view an increasing amount of one type of feedback, positive or negative, as training progresses, the feedback may serve as a confound for observed activation changes. For instance, regions responsive to positive feedback may show increasing activations over the course of training if subjects see more positive feedback in later training sessions.

Given the contrary results of the rMPFC training study and the theoretical reasons to predict that valenced feedback could interact with regulation training, we sought to uncover how positive and negative feedback differentially affect brain activity in a real-time regulation training context.

1.3.1 Neural correlates of learning from feedback

In considering how real-time feedback influences the brain, it is worth contemplating how feedback leads to learning in real-time regulation training. Theoretically, the modification of behaviour (in this case, neural/cognitive "behaviour") through positive or negative feedback represents the essence of operant conditioning (Skinner, 1953). Therefore, real-time training can be considered in light of operant conditioning principles (Weiskopf, Scharnowski et al., 2004). Operant conditioning can involve reward-based learning (where behaviours are reinforced by positive stimuli) or punishment-based learning (where inappropriate behaviours are met with negative stimuli).

A structure within the basal ganglia, the nucleus accumbens (NAc), is particularly important for reward processing. Receiving dopamine projections from the midbrain, NAc appears to help link rewarding stimuli with behavioural outcomes (Day & Carelli, 2007). Thus, if real-time training involves reward-based learning, we should expect to see NAc activation in response to positive real-time feedback. While the neural correlates of punishment-based learning are less clear, one study implicated the insula (Wächter, Lungu, Liu, Willingham, & Ashe, 2009); we might expect insula to be more active in response to negative real-time feedback.

1.4 The Present Study

The principle aim of this study was to investigate whether positive and negative feedback lead to differential patterns of brain activation in the context of real-time fMRI. The logic that regulation training may be compromised for regions activated by positive

or negative feedback, and the finding that at least one brain region (MPFC) produced contrary training results, led us to design an experiment that specifically focused on real-time feedback. While subjects view feedback which fluctuates in valence in most real-time training studies, the timing of the fluctuations cannot be controlled. In this study, which closely simulated a real-time training study, we constructed feedback in advance, allowing us to systematically manipulate its valence. The pre-constructed feedback alternated between positive feedback (moderate or strong) and negative feedback (moderate or strong), in 15, 30, or 45 second blocks.

In this simulated real-time training study, we told participants that they would be receiving real-time feedback from the parahippocampal place area (PPA), which they were to regulate up or down in 30-second blocks. The PPA is known to be activated by scene imagery (Epstein & Kanwisher, 1998); we chose it for the regulation task because we deemed it unlikely to be involved in processing positive or negative feedback. We suggested motor imagery as a down-regulation task for similar reasons. The regulation instructions were similar to those used by Weiskopf and colleagues (Weiskopf, Mathiak, et al., 2004): on up-regulation blocks, subjects were told to think of scene imagery, such as landscapes, and for down-regulation, they were told to think of motor imagery, such as playing tennis. By presenting subjects with feedback alternating in valence, we were able to compare brain activations during periods of positive and negative feedback. We expected NAc activation in response to positive feedback due to its role in reward-based learning. The MPFC was also of special interest given the reverse training finding for this region.

2 METHOD

2.1 Subjects

A total of 22 right-handed subjects (ranging from 20-31 years of age; mean age = 23.8 years; 11 female) participated in this study, recruited through our lab website. Subjects were told that the study was designed to examine their ability to regulate a particular brain region using real-time fMRI feedback. All subjects provided written consent to participate and were paid $10/hour of time spent outside the scanner, and $20/hour of time spent within the scanner.

2.1.1 Exclusion criteria

The present study required us to deceive subjects by informing them that they would be viewing real-time fMRI feedback from their own brain, when the feedback was in fact pre-constructed. We thus included a funnelled debriefing interview at the end of the study, to assess whether subjects believed this cover story throughout the scanning session (see Appendix A). Verbal responses to the interview questions were recorded by the experimenter.

The debriefing interviews were coded by the experimenter and a blind, independent rater on a scale of 1-5, with 1 indicating no suspicion, and 5 indicating strong suspicion or guessing the study's true purpose. Interrater reliability of $r = 0.80$ was achieved. We averaged the two ratings to produce one suspicion score for each subject, and used these to determine which subjects to exclude from the analysis. Six subjects with suspicion ratings above 3 were excluded in an effort to maintain the validity

of our results. While some subjects with suspicion ratings of 2 or 3 did voice doubts about the feedback, they were deemed to have sufficient credulity for their data to be trustworthy, and were included in the final analysis. One subject with a high suspicion score of 4.5 indicated that he was naïve to the feedback until a point halfway through the experiment, when he had requested a break from scanner; thus, his first three feedback runs only were included in the analysis.

Two subjects were excluded from our analysis due brain abnormalities, which also resulted in poor normalization. One additional subject was excluded due to high task-motion correlation for the PPA regulation task ($r > 0.20$). Refer to Appendix B for a summary of included and excluded subjects.

2.1.2 Included subjects

The above exclusion criteria left 13 subjects in the final analysis (age range: 20-31; mean age = 23.5 years; 7 female). For one subject, data for one session was not collected due to experimental error; another subject asked to end the experiment after four sessions; a third subject, already mentioned, had three sessions excluded due to suspicion. The remainder of subjects produced six sessions' worth of data.

2.2 Individual Difference Measures

Prior to scanning, subjects completed a simplified version of the Big Five Inventory-44 (BFI-44; John & Srivastava, 1999) on pen and paper. Following scanning, they completed a twelve-item rumination scale from the Reflection-Rumination Questionnaire (RRQ; Trapnell & Campbell, 1999) on the computer. (For a summary of

the results, refer to Appendix C.) We elected to administer the rumination measure after scanning so as to avoid raising any suspicions about the study's true purpose.

2.3 Experimental Protocol

2.3.1 Instructions

Subjects were tested individually. After signing the consent form and completing the BFI-44, they were shown the task instructions via Microsoft PowerPoint. The instructions informed subjects that we were interested in whether people could learn to control activation in the parahippocampal place area (PPA), and that their task was to regulate this region's activation up or down in 30 second blocks. They were introduced to the feedback display, which included an arrow indicating the regulation direction (red up-arrow for up-regulation, and blue down-arrow for down-regulation), and a feedback thermometer indicating the momentary activation of their PPA. Figure 2.1 illustrates the feedback display for PPA up-regulation. Subjects were told that thermometer's black midline indicated average PPA activition; thus, for up-regulation blocks, they should endeavor to keep the thermometer feedback bar above the midline, and for down-regulation blocks, they should endeavor to keep it below the midline. The fluctuating thermometer bar was blue below the mid-point, and red above it, consistent with previous real-time studies from our own lab and others (Caria et al., 2007; McCaig et al., in preparation; Rota et al., 2009; Smith, 2008).

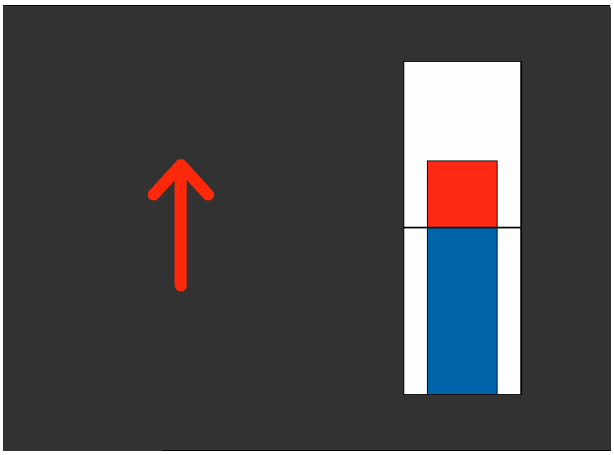

Figure 2.1: Feedback display. The arrow on the left indicated PPA regulation direction. Subjects believed that the thermometer's fluctuating feedback bar conveyed the activation level of their PPA. In reality, the feedback alternated between a positive and negative reading. The above display shows a positive feedback reading, since the thermometer bar is above the midline, corresponding to the direction of the regulation arrow.

Subjects were also (accurately) informed that the PPA is usually activated when people think about places, such as houses or landscapes. On up-regulation blocks, they were guided to visualize complete scenes, either indoors or outdoors, which may or may not be familiar to them. They were encouraged to use different strategies such as focusing on a specific scene, or switching between several scenes, to determine which strategy worked best for them. On down-regulation blocks, subjects were encouraged to imagine themselves performing a motor task, such as playing tennis, performing jumping jacks, or playing the piano. They were encouraged to adopt a first-person perspective and to focus on the sensations and motor sequences involved, but again encouraged to attempt various strategies to determine the one which was most effective.

Lastly, subjects received general instructions regarding the structure of the experiment, and the scanner environment. They were instructed to attend to and make use of the real-time feedback, but to avoid getting caught up or frustrated if the activation reading differed from their expectations. On the whole, these instructions were designed

to match those of previous real-time regulation studies as closely as possible. After advancing through the slideshow, the experimenter asked subjects to summarize the instructions, to ensure adequate comprehension.

2.3.2 Scanning protocol

Subjects were given earplugs and positioned inside the birdcage head coil with a pillow surrounding their head, for comfort and to minimize head movement. After acquiring a high-resolution in-plane structural scan and five functional dynamics (for improved co-registration), the PPA regulation task and functional scanning began. First, subjects underwent an initial regulation session with no thermometer feedback, emulating previous real-time studies from our lab.

Next, subjects performed six sessions of PPA regulation while viewing thermometer "feedback". Unbeknownst to subjects, the fluctuating bar on the thermometer was based on pre-constructed values, rather than subjects' own brain activations. Thus, while this experiment closely emulated a true real-time regulation study, we did not perform real-time signal analysis. Rather than measuring the effectiveness of viewing real-time feedback for improving PPA regulation (which has been examined by others; Weiskopf, Mathiak, et al., 2004), we were instead interested in investigating how feedback fluctuating in valence itself influences neural activation during a real-time fMRI experiment. Lastly, subjects underwent a final regulation session with no thermometer feedback.

All scanning sessions were six minutes long, beginning with the PPA up-regulation task, and alternating between up-regulation and down-regulation every 30 seconds. Between sessions, we asked subjects six questions to assess their experience

during the previous session. Subjects were asked to rate, on a scale of 1-7, the amount of energy they expended, the degree to which they used the thermometer feedback, and their perceived degree of success, for both the up- and down-regulation blocks. The questions were presented one at a time via E-Prime; subjects provided number ratings out loud, and their responses were entered by the experimenter. If a subject reported low feedback use, the experimenter encouraged him or her to pay more attention to the feedback and use it to help with the regulation task.

2.3.3 Feedback design

Six sessions of thermometer feedback were constructed in advance of the experiment (for a sample session design, see Figure 2.2). Subjects all saw the same identical feedback sessions, in counterbalanced order. To generate the feedback, we added random noise around fixed feedback values such that the thermometer bar fluctuated in a realistic fashion. Each six-minute feedback session alternated between blocks of positive and negative feedback, which were either 15, 30, or 45 seconds in length. Positive feedback meant that the thermometer bar corresponded to the desired direction for PPA regulation – i.e. above midline for up-regulation and below midline for down-regulation – while negative feedback meant that the bar fluctuated in the opposite half of the thermometer. The feedback for positive and negative blocks was either moderate (fluctuating near the thermometer midline) or strong (fluctuating toward one end of the thermometer or the other), and every session included an equal amount of each feedback level (strong positive, moderate positive, moderate negative, and strong negative). Across all six sessions, the amount of each level of feedback in PPA up-regulation blocks equaled the amount in down-regulation blocks.

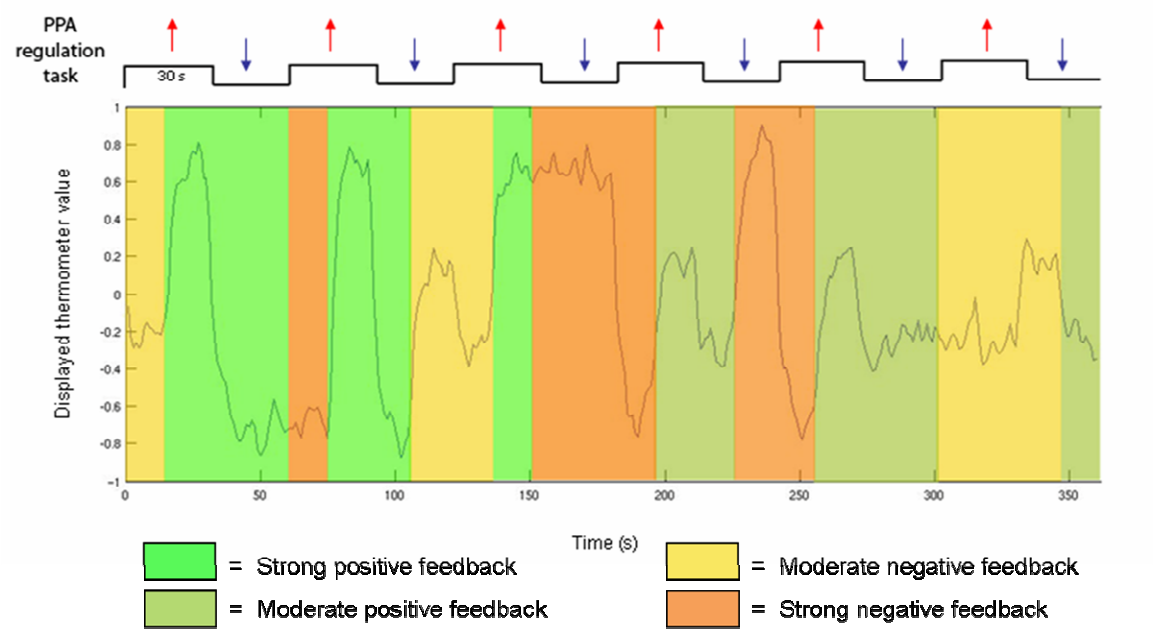

Figure 2.2: Feedback design for one sample session. Different colours indicate different levels of feedback. The red and blue arrows illustrate the PPA regulation direction, which alternated every 30 seconds. The time course illustrates the fluctuating height of the thermometer reading over the six-minute session, which depended on both the PPA task, and the type of feedback.

2.3.4 Debriefing

Following scanning, subjects were met by the experimenter and asked to complete the brief RRQ rumination scale. The experimenter then asked subjects if they could answer a few questions about their experience during scanning, leading into the debriefing interview. This six question semi-structured interview followed a funneled design such that the questions became gradually more specific, in order to gauge subjects' level of suspicion regarding the thermometer feedback's validity. At the end of the debriefing interview, all subjects were informed of the true nature of the feedback and our actual research purpose.

2.4 fMRI Data Acquisition

Imaging was conducted at the UBC MRI Research Centre, on a Philips Achieva 3.0 Tesla MRI scanner with an eight-element, six-channel phased array birdcage head coil with parallel imaging capability (SENSE; Pruessmann et al., 1999). Functional images were obtained with a T2*-weighted single shot echo-planar imaging (EPI) gradient echo sequence sensitive to fluctuations in blood-oxygen-level-dependent (BOLD) signal [time of repetition (TR) = 2000 ms; echo time (TE) = 30 ms; flip angle = 90º; field of view (FOV) = 224 × 224 × 143 mm; acquisition matrix = 80 × 80; SENSE factor = 2.0]. Each volume consisted of 36 axial slices, 6-mm thick with a 1-mm gap, acquired parallel to the anterior-commisure/posterior commisure (AC/PC) line. Each six-minute session included 180 functional volumes. Prior to functional imaging, an inversion recovery T1-weighted fast spin-echo anatomic volume was acquired, consisting of 36 3-mm axial slices with a 1-mm skip [TR = 2000 ms; TE = 10 ms; FOV = 224 × 224 × 143 mm; acquisition matrix = 240 × 235; reconstructed matrix = 480 × 480; spin echo turbo factor = 5; flip angle = 90º; inversion delay = 800 ms; SENSE factor = 2.0].

2.5 Data Analysis

2.5.1 Preprocessing

All data analysis was performed offline using SPM5 statistical parametric mapping software package (Wellcome Department of Imaging Neuroscience, London). To correct for motion, functional images were first registered to the first dynamic, and then realigned to the mean image, following a two pass procedure. Each subject's high-

resolution structural image was co-registered to the mean functional image, and then segmented to produce a grey matter image. We then normalized the segmented image to match a grey matter template, and used the same parameters to normalize the functional images. Lastly, an 8-mm full-width at half maximum (FWHM) Gaussian kernel was applied to spatially smooth the data.

2.5.2 GLM statistical analysis

We created a general linear model with five regressors: one for the PPA regulation task, and one for each level of feedback. These regressors were convolved with a canonical hemodynamic response function, and a 128 s high pass filter was applied.

We submitted individual subjects' contrast images to paired t-tests, and performed group level analyses with subject entered as a random effect. First, to assess whether subjects successfully regulated PPA activation even in the absence of valid feedback, we contrasted up-regulation with down-regulation blocks (and vice versa).

To examine the effects of feedback, we first contrasted all positive with all negative feedback. We then performed linear contrasts across the four feedback levels, in the positive and negative direction.

3 RESULTS

3.1 PPA Regulation Task Contrasts: Validation of Paradigm

To ensure that subjects were following instructions on the PPA regulation task, and that our study had sufficient power to detect activation differences, we first examined group-level activation maps for PPA up-regulation versus down-regulation. The up-regulation > down-regulation contrast revealed an activation cluster in the left PPA, which was significant at a threshold of $p < .001$ uncorrected (Figure 3.1; Table, 3.1).

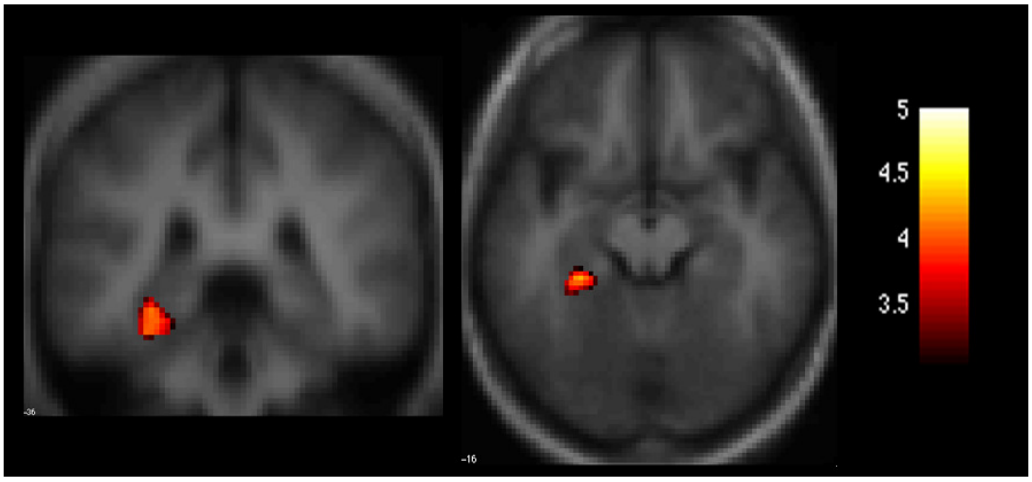

Figure 3.1: PPA up-regulation > down-regulation contrast across six feedback sessions at the group level. Activation is displayed at p < 0.005, k = 0, uncorrected, overlaid on the average of the 13 subjects' T1 structural images.

Table 3.1: PPA up-regulation > down-regulation activation table

Region	BA	MNI Coordinates			Z-value	Voxels
		x	y	z		
L Parahippocampal place area	36	-28	-34	-16	3.28	84
Height Threshold: T = 3.05, p = 0.005 Extent Threshold: k = 100 voxels						

The reverse contrast, PPA down-regulation > up-regulation, revealed numerous activations throughout the brain (Figure 3.2; Table 3.2). Most notably, extensive activation was evident in the motor cortex, particularly left motor cortex. This finding corresponded with our task instructions to imagine performing motor tasks as a method for deactivating the PPA on down-regulation blocks.

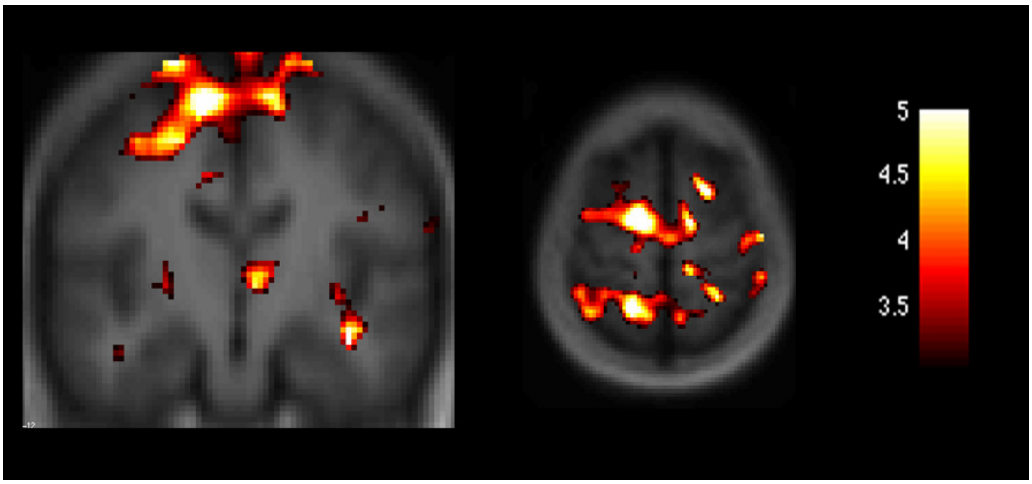

Figure 3.2: PPA down-regulation > up-regulation contrast across six feedback sessions at the group level. Activation is displayed at $p < 0.005$, $k = 0$, uncorrected, overlaid on the average of the 13 subjects' T1 structural images.

Table 3.2: PPA down-regulation > up-regulation activation table

Region	BA	MNI Coordinates			Z-value	Voxels
		x	y	z		
L Superior temporal gyrus	13	-36	-46	16	4.55	751
L Middle temporal gyrus	39	-48	-58	4	3.90	
		-64	-62	-4	3.77	
R Middle temporal gyrus		60	-48	-4	4.49	10135
R Cingulate gyrus	31	16	-30	40	4.46	
L Cingulate gyrus	24	-12	-2	36	4.45	
L Supplementary motor area	6	-14	-8	64	4.44	
R Supplementary motor area	6	12	-8	68	4.03	
L Primary motor area	4	-24	-16	52	3.80	
L Superior parietal lobule		-18	-54	64	4.07	
R Cerebellar tonsil		32	-58	-48	4.20	241
R Inferior cerebellum		14	-62	-52	2.91	
L Middle frontal gyrus	10	-32	60	28	4.14	384
L Putamen		-22	6	12	4.08	483
R Thalamus		6	-12	0	3.47	
L Anterior cerebellum		-38	-56	-28	3.96	1124
L Posterior cerebellum		-32	-60	-52	3.89	
R Superior frontal gyrus	10	32	56	20	3.91	286
R Middle frontal gyrus		38	52	16	3.48	
R Inferior frontal gyrus	45	58	18	12	3.69	392
R Middle frontal gyrus	47	44	42	-8	3.58	
Brainstem, pons		0	-26	-36	3.53	119
L Cerebellar culmen		-6	-58	-12	3.34	163
L Parahippocampal gyrus		-20	-54	-8	3.01	
L Superior temporal gyrus	22	-50	0	-4	3.21	146
L Insula		-40	6	-12	2.84	
Height Threshold: T = 3.05, p = 0.005 Extent Threshold: k = 100 voxels						

3.2 Between-Session Questions

After each scanning session, subjects were asked: (1) how much mental energy up-regulation required; (2) how much they used the feedback in order to assist them in

completing the up-regulation task; and (3) how successful they felt that they were in performing the up-regulation task. Subjects then answered the same three questions for down-regulation blocks.

Subjects' mean responses to the between-session questions are presented in Table 3.3. Paired *t*-tests revealed no significant differences between up-regulation and down-regulation responses for each question (all *p*-values > .10). We therefore averaged the up-regulation and down-regulation responses, producing single scores for energy expended, use of feedback, and perceived success for every session across all subjects.

Averaging responses across subjects, we ran linear regression analyses to determine whether responses changed across scanning sessions. Responses to questions assessing mental energy and perceived success did not change across the study (both *p*-values > .5), while reported use of feedback increased across the study, $F(1,73) = 5.66$, $\beta = .268$, $t(73) = 2.379$, $p = 0.02$. This was likely because subjects who reported low use of feedback on initial runs were encouraged by the experimenter to pay attention to and make use of the feedback on subsequent runs.

Table 3.3: Mean responses to the questions posed between regulation sessions in the scanner.

	Energy Expended		Use of Feedback		Perceived Success	
Subject	**Up-regulation**	**Down-regulation**	**Up-regulation**	**Down-regulation**	**Up-regulation**	**Down-regulation**
3	6.17	5.33	3.00	3.17	3.83	3.83
5	2.67	3.67	4.67	5.00	4.00	4.17
7	6.17	3.17	6.83	6.67	3.00	5.17
8	7.00	2.83	4.83	5.17	4.17	5.33
10	6.75	6.75	5.25	5.50	5.00	5.50
14	5.67	6.00	4.33	5.00	4.00	2.67
15	6.33	6.50	6.00	5.50	3.83	3.83
17	7.00	7.00	5.33	5.17	5.33	5.00
18	7.00	6.00	6.60	6.80	2.40	4.60
20	6.50	6.67	6.17	6.33	4.00	4.00
21	7.00	7.00	5.50	4.83	3.67	4.50
22	6.00	3.67	5.67	4.67	3.33	4.33
23	4.67	5.50	5.67	6.17	4.00	3.50
Mean	***6.07***	***5.39***	***5.37***	***5.38***	***3.89***	***4.34***
St. Dev.	*1.22*	*1.53*	*1.02*	*0.97*	*0.75*	*0.80*

Note: Ratings ranged from 1-7, with higher ratings indicating greater endorsement.

3.3 Positive Versus Negative Feedback Contrasts

The primary comparisons of interest were across the different levels of feedback. We first examined the main effects of valenced feedback by collapsing across the strong and moderate feedback blocks for positive and negative feedback, and contrasting all positive feedback with all negative feedback. Figure 3.3 and Table 3.4 display the results of the positive > negative feedback contrast. This contrast revealed an activation cluster in the nucleus accumbens extending into the left hypothalamus. In addition, activations were evident in ventral ACC and left temporal pole, both considered emotion-relevant regions. All of these clusters survived at an uncorrected threshold of $p < 0.001$.

The reverse contrast, negative > positive feedback, revealed one activation cluster in the right angular gyrus at $p < 0.005$ (extent threshold: k = 20), but no voxels from this contrast survived at $p < 0.001$.

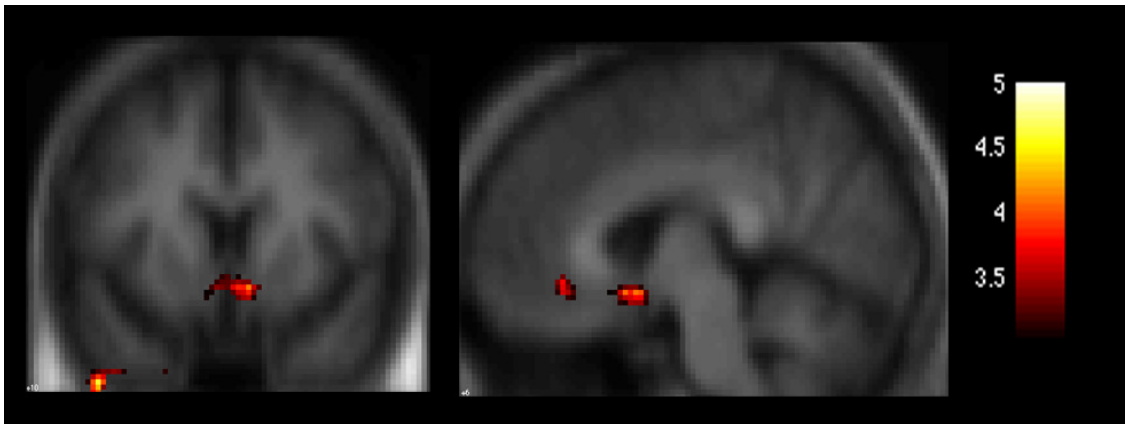

Figure 3.3 Positive feedback > negative feedback main effect contrast across six feedback sessions at the group level. Activation is displayed at p < 0.005, uncorrected, overlaid on the average of the 13 subjects' T1 structural images.

Table 3.4: Positive > negative feedback main effect activation table

Region	BA	MNI Coordinates			Z-value	Voxels
		x	y	z		
L Hypothalamus		-4	-6	-20	3.62	169
R Nucleus accumbens	25	8	10	-12	3.20	
L Parahippocampal gyrus	28	-16	-16	-24	3.19	
L Temporal pole	38	-50	10	-48	3.43	46
L Superior frontal gyrus/	9	-18	38	40	3.25	31
L Middle frontal gyrus	8	-12	26	44	3.23	66
		-16	26	56	2.92	
		-22	20	44	2.73	
L Postcentral gyrus	43	-48	-18	20	3.14	30
Ventral ACC	11	4	36	-12	3.11	25
L Precentral gyrus	9	-38	16	40	2.84	21
Height Threshold: T = 3.05, $p = 0.005$ Extent Threshold: k = 20 voxels						

To test whether the effects of feedback valence were enhanced at stronger levels of feedback, we performed a linear contrast, giving stronger contrast weights to strong

than moderate feedback. The positive linear contrast showed small activations clusters in left insula, and right and left postcentral gyri, but surprisingly, none of the regions observed in the positive > negative main effect contrast were active at $p < 0.005$, k = 20. The negative linear contrast revealed no significant activations at this threshold.

3.3.1 Inter-subject variation in positive and negative feedback activations

To examine whether the null results observed for the negative > positive feedback contrasts resulted from greater variability between subjects' responses to negative feedback, we created standard deviation images for strong negative and strong positive feedback. These images were created by computing the standard deviation of subjects' activation beta weights for every voxel, and displaying those voxels exceeding standard deviations cutoffs of 1.5 and 2. To assess the effects of positive and negative feedback independently, we used the strong positive or strong negative regressor against a baseline which collapsed across the two moderate feedback blocks. That is, we looked at the variation in activation for strong negative feedback versus both levels of moderate feedback, and compared this with the variation in activation for strong positive feedback versus both levels of moderate feedback. As can be seen in Figure 3.4, neural responses to strong negative feedback were not more variable than to positive feedback.

For both positive and negative feedback, high variability was present in bilateral insula. We also note that the MPFC showed beta weight standard deviations > 1.5 in response to strong positive feedback.

Strong positive feedback > combined moderate feedback baseline:

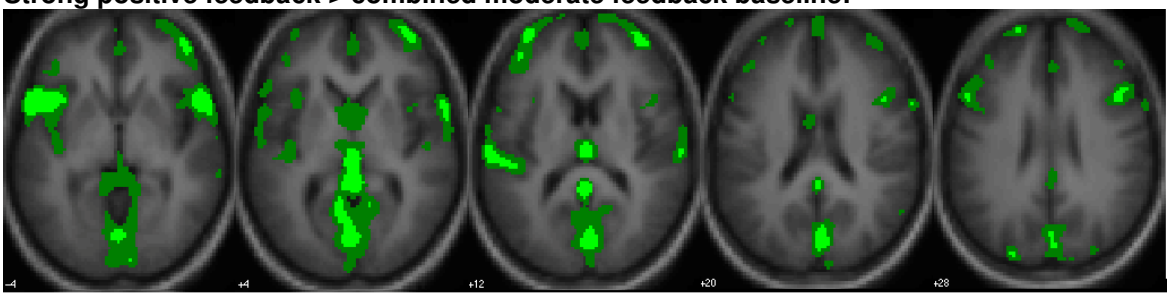

Strong negative feedback > combined moderate feedback baseline:

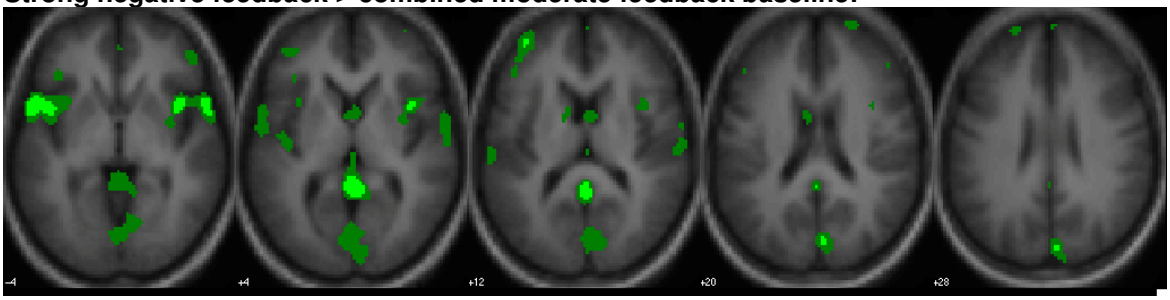

Figure 3.4: Standard deviation images for strong positive > combined moderate feedback contrast (above) and strong negative > combined moderate feedback contrast (below). Darker green represents regions with standard deviations > 1.5; brighter green represents regions with standard deviations > 2.0.

3.4 Regional Activations by Feedback Level

Because the linear contrast results for positive feedback did not resemble the main effect contrast for positive > negative feedback, we examined the activation levels across each feedback level within regions of particular interest. After defining each region, we created a single regressor for each feedback level (strong positive, moderate positive, moderate negative, and strong negative), and computed voxel-wise beta weights from these regressors for every subject. Averaging across voxels produced one beta weight for each feedback level per subject; we then centred individual subjects' beta weights about their own mean activations before computing an overall mean beta weight for each condition.

3.4.1 Nucleus accumbens (NAc) results

Given its involvement in reward and implicit learning, we first examined activations within the NAc. We defined this cluster using a sphere of 2-mm radius around the peak activation at [8 10 -12] observed in the positive > negative feedback contrast. Figure 3.5 illustrates the results as a bar graph. Unexpectedly, although both positive feedback levels showed higher NAc activation than both negative feedback levels, the effects were most pronounced for moderate positive and negative feedback.

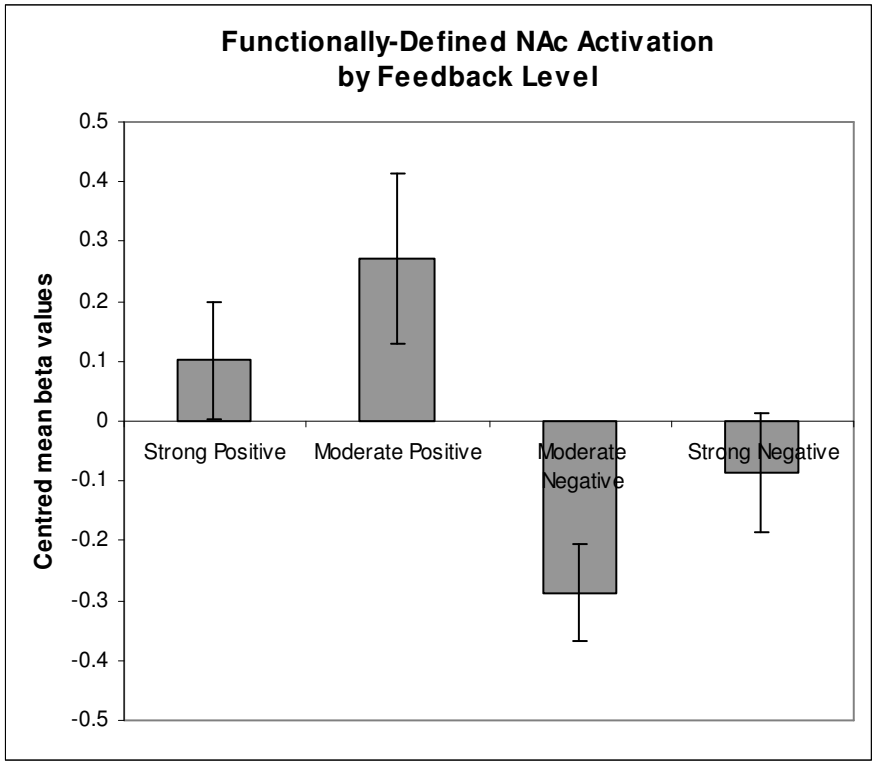

Figure 3.5: NAc activation across four levels of feedback. At each feedback level, individual subjects' activation beta values were centred, and an average centred beta value was then computed across subjects.

3.4.2 Ventral anterior cingulate cortex (vACC) results

The vACC is an affective region which was also observed in the positive > negative feedback contrast. We defined a cluster around the peak vACC activation at [4

36 -12] using a 2-mm radius sphere. Similar to the results for NAc, we found that the vACC effects were stronger at moderate levels of feedback (Figure 3.6).

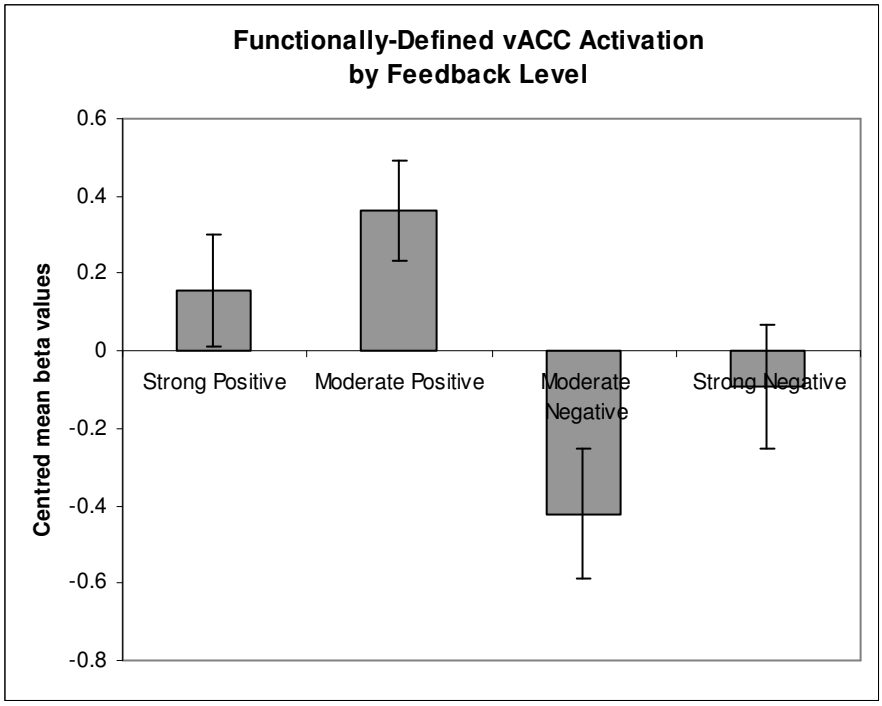

Figure 3.6: vACC activaton across four levels of feedback. At each feedback level, individual subjects' activation beta values were centred, and an average centred beta value was then computed across subjects.

3.4.3 Exploratory positive versus negative feedback contrast for moderate feedback

Given the unexpected activation pattern observed in NAc and vACC, where the strongest effects were observed for moderate levels of feedback, we performed unplanned exploratory contrasts between moderate positive and moderate negative levels of feedback. The results of the moderate positive > moderate negative feedback contrast are presented in Figure 3.7. In addition to NAc and vACC, this contrast revealed activation in the ventomedial prefrontal cortex (vMPFC). The reverse contrast again produced one activation cluster in the right angular gyrus at $p < 0.005$, k = 20.

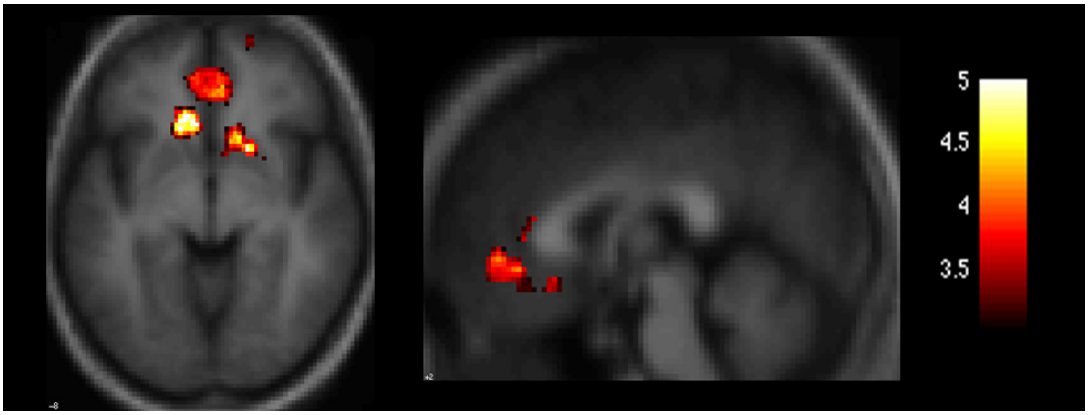

Figure 3.7: Exploratory contrast for moderate positive feedback > moderate negative feedback contrast across six feedback sessions at the group level. Activation is displayed at $p < 0.005$, $k = 0$, uncorrected, overlaid on the average of the 13 subjects' T1 structural images.

3.4.4 Medial prefrontal cortex (MPFC) results

Given our lab's previous findings showing a performance decline over real-time feedback training of MPFC, the MPFC was of particular a priori interest. The activation observed in the above moderate positive > moderate negative feedback contrast was more ventral to the one identified in the real-time training study; thus, we identified MPFC regions according to both functional criteria from the feedback contrast and previously published structural criteria.

We functionally specified vMPFC by creating a 2-mm sphere around the peak vMPFC activation at [6 42 -8] from the previous contrast. Figure 3.8 presents the results. We found that for vMPFC, both levels of positive feedback produced higher activations than moderate negative feedback; however, the activation for strong negative feedback was also higher than for moderate negative feedback, and approximately equivalent to the positive feedback activations.

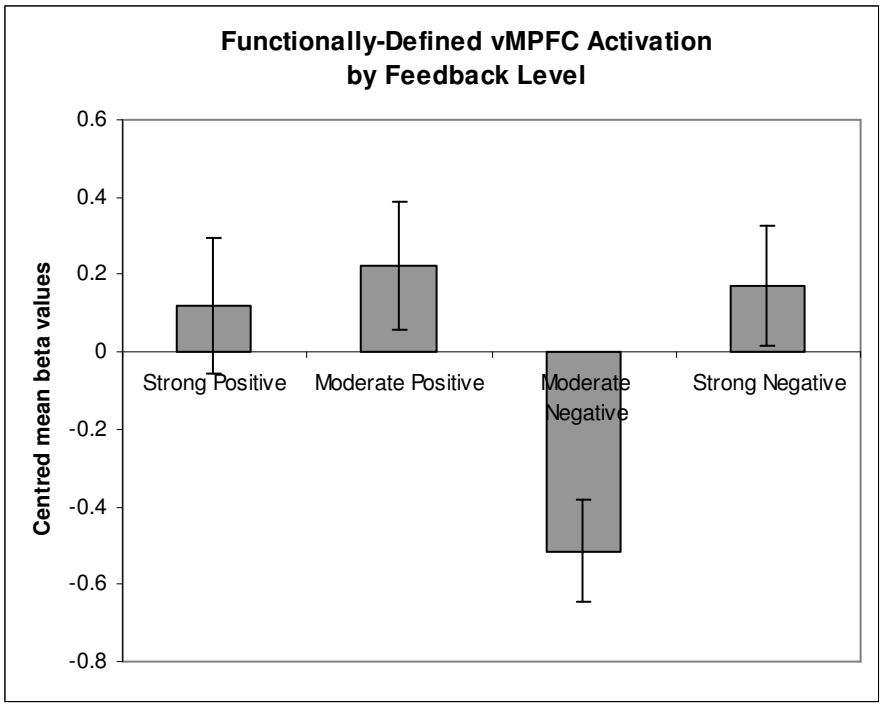

Figure 3.8: vMPFC activation (functionally-defined) across four levels of feedback. At each feedback level, individual subjects' activation beta values were centred, and an average centred beta value was then computed across subjects. The vMPFC regions was based on the contrast for moderate positive > moderate negative feedback.

We structurally defined a region within MPFC using WFU PickAtlas Version 2.1 software (Maldjian, Laurienti, Kraft, & Burdette, 2003). By intersecting the mask for medial frontal gyrus with Brodmann Area 10, with upper and lower cutoffs of z = +30 and +10, we delineated a region which closely resembled that from our lab's MPFC regulation study (Smith, 2008).

Figure 3.9 presents the results, which closely resemble the vMPFC results above. Again, strong positive, moderate positive, and strong negative feedback show similar levels of activation, while moderate negative feedback was deactivated relative to the other three feedback levels.

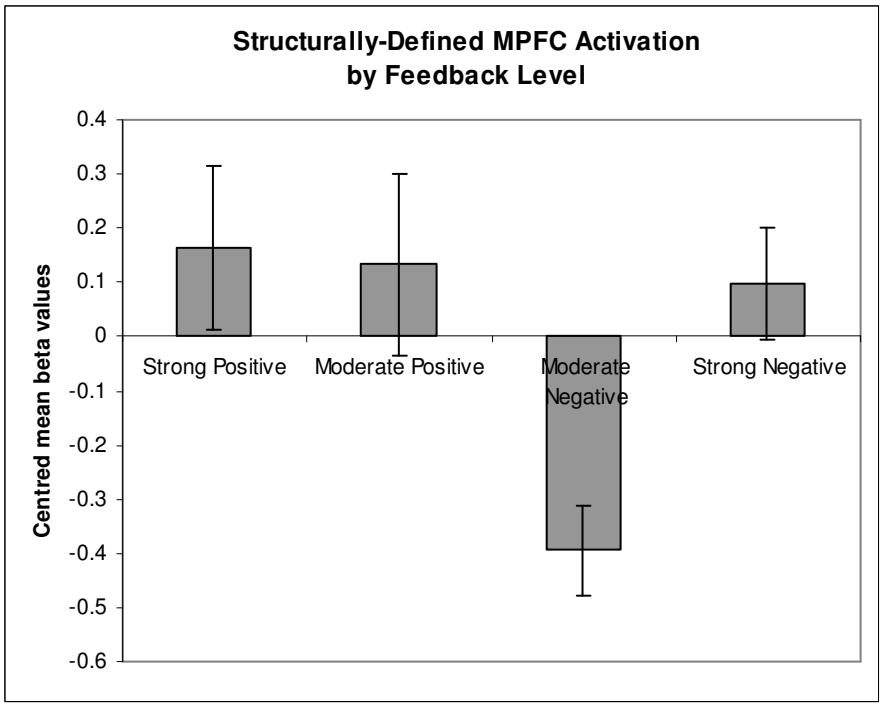

Figure 3.9: MPFC activaton (structurally-defined) across four levels of feedback. At each feedback level, individual subjects' activation beta values were centred, and an average centred beta value was then computed across subjects. The MPFC was defined based on our lab's previous real-time investigation of MPFC regulation (Smith, 2008).

3.5 Regression Analyses with Individual Difference Measures

To examine whether activations for positive or negative feedback were related to individual difference measures, we ran exploratory regression analyses examining the relationship between each personality scale (extraversion, agreeableness, conscientiousness, neuroticism, and openness to experience) and feedback valence contrasts. We used the same contrasts from Section 3.3.1, separately contrasting strong positive and strong negative feedback against a moderate feedback baseline. Given the number of regression analyses and their exploratory nature, we raised the height threshold to an FDR-corrected $p = 0.05$.

Running second level random effects analysis with personality scale scores entered as a covariate, one regression analysis produced significant activations above the given height threshold. Extraversion scores significantly predicted activations in response to strong negative feedback (see Figure 3.10). Notably, the predicted activations included a cluster in rMPFC.

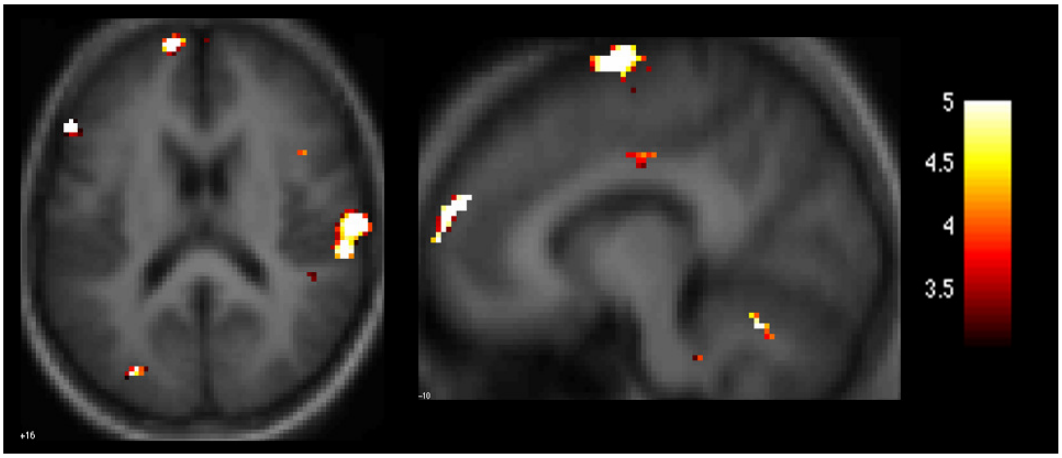

Figure 3.10: Exploratory regression analysis constrasting strong negative feedback against both levels of moderate feedback, with extraversion scores entered as a covariate. Activations are displayed at $p < 0.001$, $k = 0$.

3.6 Funnelled Debriefing Interview

In the debriefing interviews (see Appendix A), which were delivered before informing subjects of the true purpose of the study, subjects reported a range of experiences during the scanning. Some indicated that the feedback was helpful, while others said it was distracting (some said it was both helpful and distracting). Similarly, reports were mixed about whether subjects preferred performing the regulation task with or without the feedback.

Most subjects reported trying out various strategies on the PPA regulation task. Usually, they described specific types of mental imagery which they felt had helped them

up-regulate or down-regulate the PPA. For instance, one subject reported that on up-regulation blocks, she initially thought about personal experiences of scenes, but later switched to thinking about pictures such as her computer's desktop background. Another subject reported on down-regulation blocks thinking of running, math equations, counting and playing guitar. Often, subjects indicated certain strategies working better than others. Given that the feedback did not reflect subjects' actual thought processes and was negative half the time, subjects' attempts to find an appropriate strategy illustrated that they were putting effort into the task and making use the feedback.

The fourth interview question asked subjects directly, "Can you tell me what this experiment was about?" Responses on this question were used to assess whether subjects still believed the cover story from the beginning of the experiment. To subtly probe for suspicion, we next asked whether subjects had noticed a link between their thoughts and the feedback. Lastly, subjects were asked directly whether they had any suspicions regarding the study's true purpose. Responses to these three questions, which varied widely, were later considered by two raters who made holistic judgments regarding subjects' suspicion levels during the experiment. Raters additionally considered spontaneous comments indicating suspicion toward the feedback made at any point during the debriefing interview. Responses such as "The feedback seemed related to what I was thinking for the most part, but at times seemed random or directly opposite to what I expected" were taken to indicate that the subject was naive to the true nature of the feedback. However, if subjects indicated that the experiment was about something other than PPA regulation, or that the feedback was not tied to their thoughts, they were generally assigned high suspicion scores.

Coding the debriefing interviews along a 5-point scale ranging from 1 (not suspicious) to 5 (highly suspicious), the two raters achieved inter-rater reliability of r = 0.80, and their ratings were averaged to produce one suspicion score per subject.

For the initial group of 22 subjects, the mean suspicion score was 2.72. For the final group of 13 subjects (not counting one subject whose high suspicion score reflected later runs which were excluded from his first level analysis), the mean suspicion score was 1.75 and the median score was 1.5 (SD = 0.75).

Thus, on the whole, subjects varied significantly in terms their acceptance of the experimental cover story, but our final sample included subjects who were quite credulous, bolstering our confidence in their neural responses to the feedback.

4 DISCUSSION

This study set out to explore whether positive and negative feedback produced different activation patterns in a proxy real-time fMRI regulation study. We found that positive feedback yielded consistent activations in reward centres, including nucleus accumbens (NAc), and emotion centres, including ventral anterior cingulate cortex (vACC) and left temporal pole. By contrast, negative feedback led only to activation of right angular gyrus over positive feedback. On the whole, few cortical activations were observed for either contrast, a finding with optimistic implications for training paradigms.

The NAc activation confirmed that positive feedback was rewarding. However, when we examined the activation pattern in NAc across the four feedback levels, the results were surprising. Instead of observing the greatest activation for strong positive feedback, and the least activation for strong negative feedback, we observed the greatest activation for *moderate* positive feedback and the least activation for *moderate* negative feedback. This finding prompts further consideration of the role of NAc in reward processing.

A seminal study on non-human primates showed that rather than firing in linear relationship with the strength of reward, dopamine neurons (which have primary NAc projections) respond to reward uncertainty. Using a conditioning paradigm, Fiorillo and colleagues manipulated the probability that a reward would be delivered two seconds after a visual cue (Fiorillo, Tobler, & Schultz, 2003). The cue indicated the likelihood of reward, which ranged from $P = 0$ to $P = 1.0$. These researchers found that sustained activation of dopamine neurons was related to the uncertainty of an upcoming reward, with the greatest change in activity occurring at peak predictive uncertainty ($P = 0.5$).

Neuroimaging research on humans has produced convergent results (Berns, McClure, Pagnoni, & Montague, 2001). When subjects were given water or juice rewards according to a predictable or unpredictable schedule, NAc activation was greatest for unpredictable rewards. Furthermore, subjects' preferences for juice or water did not predict activation in reward regions.

In our study, moderate positive feedback fluctuating near the thermometer midline may have promoted uncertainty about whether the feedback would remain rewarding. If it did, the NAc likely showed the greatest activation in response to moderate positive feedback based on the increased uncertainty associated with it. In the case of negative feedback, research has found dopamine firing suppression (using single cell recording; Fiorillo et al., 2003) and reward centre deactivation (using fMRI; Pagnoni, Zink, Montague, & Berns, 2002) when expected rewards are not delivered. However, this does not fully explain why in the present experiment, moderate negative feedback led to greater NAc deactivation than strong negative feedback. We speculate that uncertainty may exacerbate the effects of negative feedback as well.

The lack of activations for the negative versus positive feedback contrast was also somewhat surprising. We consider it unlikely that low power was to blame, since we detected activations for positive versus negative feedback and for PPA up- and down-regulation at the same threshold. From subjects' debriefing reports, we gathered that different subjects responded to the negative feedback in different ways: certain subjects said that negative feedback motivated them to try harder on the task, other subjects found it alarming, and still others were discouraged by it. Yet when we compared the inter-subject variation maps for strong negative and strong positive feedback, there appeared to

be more between-subject variability for positive feedback than negative feedback. Interestingly, however, both positive and negative variation maps showed high variation in bilateral insula.

The insula has been linked to punishment learning, the process of learning to avoid responses associated with negative stimuli (Wächter, Lungu, Liu, Willingham, & Ashe, 2009). While conceptually, punishment is the complement of reward learning, it may not have the same lasting effects; B. F. Skinner himself questioned the effectiveness of negative reinforcement (punishment) for long-term retention (Skinner, 1953). In a procedural sequence learning experiment comparing reward versus punishment, Wächter and colleagues (2009) found that while punishing negative responses produced short-term behavioural effects, only reward-based training led to lasting changes learning-based over a control condition.

The questionable efficacy of punishment learning, along with subjects' variegated reports of their responses to negative real-time feedback, raises the possibility that negative feedback may be less useful than positive feedback for improving brain regulation with real-time fMRI. Bray and colleagues successfully trained people to activate localized brain regions using a reward-based shaping procedure (Bray, Shimojo, & O'Doherty, 2007). Instead of presenting live visual feedback, they offered subjects monetary rewards for activating a cortical finger or toe region beyond a certain threshold, which increased based on subjects' prior successes. Money was never withdrawn for unsuccessful performance. In deCharms and colleagues' study examining regulation of rACC and corresponding pain experience, subjects viewed a fire, which increased or decreased in size according to rACC activation (deCharms et al., 2005). (A fluctuating

line graph of the activation was also presented.) Both healthy subjects and pain patients not only improved at regulating rACC, but also reported significant changes in their pain experience. Based on these findings, it may be worth further exploring methods of feedback that capitalize on reward-based reinforcement learning.

4.1 Real-time Training of Emotion Regions

Emotion regions represent appealing targets for real-time training due to the potential clinical applications of emotion regulation gains. In the present study, we were particularly interested in the effects of real-time feedback on MPFC, given our lab's contrary training findings for this region (Smith, 2008). MPFC activation was not observed for the main feedback contrasts, but in an exploratory contrast comparing moderate positive feedback to moderate negative feedback, we found activation a ventral portion of MPFC. The specific activation pattern in this region was similar to the pattern observed when we structurally defined a more dorsal region of MPFC: strong positive, moderate positive, and strong negative feedback all produced greater activations than moderate negative feedback.

Though this activation pattern may appear unusual, previous studies have implicated MPFC in both positive and negative emotional experiences. For instance, processing positive traits about one's self has been shown to activate this region (Fossati et al., 2003; Moran, Macrae, Heatherton, Wyland, & Kelley, 2006). However, watching videos of personal failures led to MPFC activation in athletes (Davis IV et al., 2008). Our MPFC activation findings are consistent with the notion that this region is involved

in processing both positive and strong negative emotions, which may account for difficulties in training MPFC regulation using valenced feedback.

Our positive versus negative feedback contrast pointed to vACC and left temporal pole, two other emotion-relevant regions. We speculate that the direction of regulation training (up-regulation or down-regulation) might matter for emotion regions. Specifically, if positive feedback activates such regions, up-regulation may be facilitated and down-regulation impaired. Interestingly, of the three real-time studies published to date targeting emotion regions, two reported activation gains for up-regulation only. Caria and colleagues (2007) showed that subjects learned to up-regulate the right anterior insula (which was structurally defined) after feedback training. Johnston and colleagues (2010) reported similar results for a functionally-defined emotion network. Neither research group reported down-regulation of the target regions; it is unclear whether this absence reflects failed efforts at down-regulation training or a lack of focus on down-regulation altogether.

Another study did investigate down-regulation for subgenual ACC, an affective region implicated in major depression (Hamilton, Glover, Hsu, Johnson, & Gotlib, 2010). Subjects were instructed to alternate between down-regulating this region and rest, over an initial pre-training session, two real-time feedback sessions, and a post-training session. Although activation decreases on down-regulation blocks were reported for the two feedback sessions, the decrease diminished on the second feedback session relative to the first. Furthermore, the difference between the second feedback session and the initial no-feedback session was only marginally significant, and a post-training regulation session yielded no training effects. Thus, the results of this study are far from conclusive

regarding the feasibility of improving down-regulation of emotion region using real-time feedback.

4.2 Experimental Design Limitations

While the present experiment did detect activations at an uncorrected height threshold of $p = 0.001$, none of the feedback valence activations survived an FDR correction. Although an effort was made to equalize the amount of positive and negative feedback that occurred during up-regulation blocks and down-regulation blocks over the full experiment, and although each session contained an equal balance of positive and negative feedback, within each session, the type of feedback covaried with PPA regulation condition. Certain sessions contained more positive feedback (and less negative feedback) in the up-regulation condition, while other sessions contained more positive feedback (and less negative feedback) in the down-regulation condition. Since the proportion of positive and negative feedback was balanced across the whole experiment, this design issue was not an experimental confound, but it may have limited our power to detect activations resulting from the feedback valence manipulation.

4.2.1 Lack of baseline condition

It is important to note that because this study did not include a baseline condition, we cannot make inferences about the isolated effects of positive or negative real-time feedback; we can only compare one to the other. If certain regions are activated in response to *both* positive and negative feedback, they could not be identified by the present experiment. We elected not to include a baseline condition for two reasons.

First, we felt that if we used "neutral" feedback as a baseline, the feedback might not be interpreted as genuinely neutral. Feedback hovering around the midline on the thermometer might be interpreted as negative rather than neutral, as subjects might interpret it as failing to rest in the desired half of the thermometer. Second, we designed the experiment to emulate true real-time regulation studies as closely as possible, and thus, presented feedback continuously over the training runs.

While the lack of baseline condition limits our ability to draw inferences about the isolated effects of positive and negative real-time feedback, this does not pose a serious problem for our research aim. In real-time regulation experiments, feedback is presented continuously throughout the training runs, and thus, if certain regions are responsive to both positive and negative feedback, they should be active throughout the training runs, and should not *interact* with the feedback.

However, our paradigm does raise the possibility that the activations observed in the positive versus negative feedback contrast reflect *deactivations* caused by negative feedback, as opposed to activations caused by positive feedback. For instance, the vACC has been consistently deactivated by cognitively demanding tasks (Bush, Luu, & Posner, 2000). It is possible that the apparent vACC activation in the positive versus negative feedback contrast resulted from subjects exerting more cognitive effort during negative feedback periods, thereby deactivating the region.

4.3 Comparability with Real-Time Regulation Studies

We designed this study to be as comparable as possible with true real-time experiments. To make the feedback appear realistic, we added random noise around the four thermometer values. Every session included an even balance of positive and negative feedback; that is, we did not build improvement into the study by increasing the amount of positive feedback in later runs. In true real-time regulation studies showing positive training effects, many subjects see greater amounts of positive feedback as they learn to use the feedback. However, in most of these studies, a subset of subjects showed the opposite pattern, becoming worse at the regulation task, and thus seeing more negative feedback as the study progresses.

The even balance of positive and negative feedback, and the lack of connection between subjects' thoughts and the thermometer reading, raised suspicions about the true nature of the feedback among certain subjects. From the debriefing questionnaire, it was apparent that subjects suspicion levels ranged from complete naiveté to declaring that the feedback was not real, with many subjects falling somewhere in between (for example, asking whether all of the feedback had been accurate). However, we note that in previous training studies from our lab, some subjects who saw true real-time feedback reflecting their own brain activations also questioned its validity. For this reason, we excluded subjects with high suspicion, but felt that subjects with low to moderate suspicion were sufficiently naive so as to respond to the feedback in a realistic manner, and thus included them in our final sample.

4.4 Conclusions

This study focused on the mechanism of real-time regulation training – namely, the feedback itself. The small number of cortical activations observed in our feedback comparisons may be good news for real-time training, suggesting that feedback should not interact with regulation for most cortical regions. In real-time studies, positive feedback likely facilitates reward-based learning; it is less clear whether negative feedback enhances learning. Future training efforts may benefit from developing feedback methods that emphasizes positive rewards, for instance, using a shaping procedure to reward successively greater improvements (Bray et al., 2007).

While emotion-relevant regions represent some of the most alluring targets for real-time studies, they may also be most sensitive to feedback fluctuating in valence. New training strategies, such as separating feedback from regulation periods, may be necessary in cases where feedback interacts with the regulation task. On the whole, real-time regulation training is still a developing area, with much room for future work. Both positive and negative findings should be published in order for our understanding to advance. Whether or not real-time regulation training will one day gain the status of an empirically supported treatment remains to be seen.

REFERENCES

Bagarinao, E., Nakai, T., & Tanaka, Y. (2006). Real-time functional MRI: Development and emerging applications. *Magnetic Resonance in Medical Sciences : MRMS : An Official Journal of Japan Society of Magnetic Resonance in Medicine, 5*(3), 157-65.

Beauregard, M., & Levesque, J. (2006). Functional magnetic resonance imaging investigation of the effects of neurofeedback training on the neural bases of selective attention and response inhibition in children with attention-deficit/hyperactivity disorder. *Applied Psychophysiology and Biofeedback, 31*(1), 3-20.

Berns, G. S., McClure, S. M., Pagnoni, G., & Montague, P. R. (2001). Predictability modulates human brain response to reward. *Journal of Neuroscience, 21*(8), 2793.

Bray, S., Shimojo, S., & O'Doherty, J. P. (2007). Direct instrumental conditioning of neural activity using functional magnetic resonance imaging-derived reward feedback. *Journal of Neuroscience, 27*(28), 7498.

Bush, G., Luu, P., & Posner, M. I. (2000). Cognitive and emotional influences in anterior cingulate cortex. *Trends in Cognitive Sciences, 4*(6), 215-222.

Cannon, R., Lubar, J., Congedo, M., Thornton, K., Towler, K., & Hutchens, T. (2007). The effects of neurofeedback training in the cognitive division of the anterior cingulate gyrus. *The International Journal of Neuroscience, 117*(3), 337-57.

Caria, A., Veit, R., Sitaram, R., Lotze, M., Weiskopf, N., Grodd, W., et al. (2007). Regulation of anterior insular cortex activity using real-time fMRI. *NeuroImage, 35*(3), 1238-46.

Christopher deCharms, R. (2008). Applications of real-time fMRI. *Nature Reviews Neuroscience, 9*(9), 720-729.

Critchley, H. D., Melmed, R. N., Featherstone, E., Mathias, C. J., & Dolan, R. J. (2001). Brain activity during biofeedback relaxation: A functional neuroimaging investigation. *Brain : A Journal of Neurology, 124*(Pt 5), 1003-12.

Davis IV, H., Liotti, M., Ngan, E., Woodward, T. S., Van Snellenberg, J. X., van Anders, S. M., et al. (2008). fMRI BOLD signal changes in elite swimmers while viewing videos of personal failure. *Brain Imaging and Behavior, 2*, 84-93.

Day, J. J., & Carelli, R. M. (2007). The nucleus accumbens and pavlovian reward learning. *The Neuroscientist, 13*(2), 148.

decharms, R. C. (2007). Reading and controlling human brain activation using real-time functional magnetic resonance imaging. *Trends in Cognitive Sciences, 11*(11), 473-481.

deCharms, R. C. (2008). Applications of real-time fMRI. *Nature Reviews Neuroscience, 9*, 720-729.

deCharms, R. C., Christoff, K., Glover, G. H., Pauly, J. M., Whitfield, S., & Gabrieli, J. D. (2004). Learned regulation of spatially localized brain activation using real-time fMRI. *NeuroImage, 21*(1), 436-43.

deCharms, R. C., Maeda, F., Glover, G. H., Ludlow, D., Pauly, J. M., Soneji, D., et al. (2005). Control over brain activation and pain learned by using real-time functional MRI. *Proceedings of the National Academy of Sciences of the United States of America, 102*(51), 18626-31.

Epstein, R., & Kanwisher, N. (1998). A cortical representation of the local visual environment. *Nature, 392*(6676), 598-601.

Farb, N. A. S., Segal, Z. V., Mayberg, H., Bean, J., McKeon, D., Fatima, Z., et al. (2007; in press). Attending to the present: Mindfulness meditation reveals distinct neural modes of self-reference. *SCAN,*

Fiorillo, C. D., Tobler, P. N., & Schultz, W. (2003). Discrete coding of reward probability and uncertainty by dopamine neurons. *Science, 299*(5614), 1898.

Fossati, P., Hevenor, S. J., Graham, S. J., Grady, C., Keightley, M., Craik, F., et al. (2003). In search of the emotional self: An fMRI study using positive and negative emotional words. *American Journal of Psychiatry, 160*(11), 1938-1945.

Hamilton, J. P., Glover, G. H., Hsu, J. J., Johnson, R. F., & Gotlib, I. H. (2010) Modulation of subgenual anterior cingulate cortex activity with real-time neurofeedback. *Human Brain Mapping.*

John, O. P., & Srivastava, S. (1999). The Big 5 trait taxonomy: History, measurement, and theoretical perspectives. In L. A. Pervin&O. P. John (Eds.), *Handbook of personality: Theory and research* (2nd ed., pp. 102-138). New York: Guilford.

Johnston, S., Boehm, S., Healy, D., Goebel, R., & Linden, D. (2010). Neurofeedback: A promising tool for the self-regulation of emotion networks. *NeuroImage, 49*(1), 1066-1072.

Keramatian, K. (2009). *Modulation of Rostrolateral Prefrontal Cortex during Cognitive Introspection using Real-Time fMRI.* Master's thesis. Vancouver: University of British Columbia.

Kiehl, K. A. (2006). A cognitive neuroscience perspective on psychopathy: Evidence for paralimbic system dysfunction. *Psychiatry Research, 142*(2-3), 107-128.

Laconte, S. M., Peltier, S. J., & Hu, X. P. (2007). Real-time fMRI using brain-state classification. *Human Brain Mapping, 28*, 1033-1044.

Lane, R. D., Fink, G. R., & Chau, P. ML., & Dolan, RJ (1997). neural activation during selective attention to subjective emotional responses. *Neuroreport, 8*, 3969-3972.

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

McCaig, R. G., Dixon, M., Keramatian, K., Liu, I., Christoff, K. (in preparation). Improved modulation of rostrolateral prefrontal cortex using real-time fMRI training and meta-cognitive awareness.

Moran, J. M., Macrae, C. N., Heatherton, T. F., Wyland, C. L., & Kelley, W. M. (2006). Neuroanatomical evidence for distinct cognitive and affective components of self. *Journal of Cognitive Neuroscience, 18*(9), 1586-1594.

Moss, D., & Kirk, L. (2004). Evidence-based practice in biofeedback and neurofeedback. Foreword to C. Yucha & C. Gilbert, *Evidence-based practice for biofeedback assisted behavioral therapy* (pp. iii–ix). Wheat Ridge, CO: Association for Applied Psychophysiology and Biofeedback.

Myers, C. E., Shohamy, D., Gluck, M. A., Grossman, S., Kluger, A., Ferris, S., et al. (2003). Dissociating hippocampal versus basal ganglia contributions to learning and transfer. *Journal of Cognitive Neuroscience, 15*(2), 185-193.

Ochsner, K. N., Ray, R. D., Cooper, J. C., Robertson, E. R., Chopra, S., Gabrieli, J. D. E., et al. (2004). For better or for worse: Neural systems supporting the cognitive down- and up-regulation of negative emotion. *NeuroImage, 23*(2), 483-499.

Oschner, K. N., Ray, R. D., Cooper, J. C., Robertson, E. R., Chopra, S., Gabrieli, J. D. E., et al. (2004). For better or for worse: Neural systems supporting the cognitive down- and up-regulation of negative emotion. *NeuroImage, 23*, 483-499.

Owen, A. M., Martin, R. C., Boly, M., Davis, M. H., Laureys, S., & Pickard, J. D. (2006). Detecting awareness in the vegetative state. *Science, 313*, 1402.

Pagnoni, G., Zink, C. F., Montague, P. R., & Berns, G. S. (2002). Activity in human ventral striatum locked to errors of reward prediction. *Nature Neuroscience, 5*(2), 97-98.

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

Posse, S., Fitzgerald, D., Gao, K., Habel, U., Rosenberg, D., Moore, G. J., et al. (2003). Real-time fMRI of temporolimbic regions detects amygdala activation during single-trial self-induced sadness. *NeuroImage, 18*(3), 760-8.

Rota, G., Sitaram, R., Veit, R., Erb, M., Weiskopf, N., Dogil, G., et al. (2009). Self-regulation of regional cortical activity using real-time fMRI: The right inferior frontal gyrus and linguistic processing. *Human Brain Mapping, 30*(5), 1605-1614.

Santos, J. M., Wright, G. A., & Pauly, J. M. (2004). Flexible real-time magnetic resonance imaging framework. *Conference Proceedings : ... Annual International Conference of the IEEE Engineering in Medicine and Biology Society. IEEE Engineering in Medicine and Biology Society. Conference, 2*, 1048-51.

Schneider, F., Backes, V., & Mathiak, K. (2009). Brain imaging: On the way toward a therapeutic discipline. *European Archives of Psychiatry and Clinical Neuroscience, 259*, 143-147.

Shanks, D. R., & St. John, M. F. (2010). Characteristics of dissociable human learning systems. *Behavioral and Brain Sciences, 17*(03), 367-395.

Shanks, D. R., & St. John, M. F. (2010). Characteristics of dissociable human learning systems. *Behavioral and Brain Sciences, 17*(03), 367-395.

Smith, R. M. (2008). *Regulating the Anterior Medial Prefrontal Cortex: Exploratory Investigation of Real-time Training.* Master's thesis. Vancouver: University of British Columbia.

Sitaram, R., Caria, A., Veit, R., Gaber, T., Rota, G., Kuebler, A., et al. (2007). fMRI brain-computer interface: A tool for neuroscientific research and treatment. *Computational Intelligence and Neuroscience, 2007*, 1.

Sorger, B., Dahmen, B., Reithler, J., Gosseries, O., Maudoux, A., Laureys, S., et al. (2009). Another kind of BOLD response: Answering multiple-choice questions via online decoded single-trial brain signals. *Progress in Brain Research,* , 275-292.

Skinner, B. F. (1953). *Science and human behavior*. Oxford, England: Macmillan.

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

Trapnell, P. D., & Campbell, J. D. (1999). Private self-consciousness and the five-factor model of personality: Distinguishing rumination from reflection. *Journal of Personality and Social Psychology, 76*(2), 284-304.

Wächter, T., Lungu, O. V., Liu, T., Willingham, D. T., & Ashe, J. (2009). Differential effect of reward and punishment on procedural learning. *Journal of Neuroscience, 29*(2), 436.

Weiskopf, N., Mathiak, K., Bock, S. W., Scharnowski, F., Veit, R., Grodd, W., et al. (2004). Principles of a brain-computer interface(BCI) based on real-time functional magnetic resonance imaging(fMRI). *IEEE Transactions on Biomedical Engineering, 51*(6), 966-970.

Weiskopf, N., Scharnowski, F., Veit, R., Goebel, R., Birbaumer, N., & Mathiak, K. (2004). Self-regulation of local brain activity using real-time functional magnetic resonance imaging (fMRI). *Journal of Physiology-Paris, 98*(4-6), 357-373.

Weiskopf, N., Sitaram, R., Josephs, O., Veit, R., Scharnowski, F., Goebel, R., et al. (2007). Real-time functional magnetic resonance imaging: Methods and applications. *Magnetic Resonance Imaging, 25*(6), 989-1003.

Weiskopf, N., Mathiak, K., Bock, S. W., Scharnowski, F., Veit, R., Grodd, W., et al. (2004). Principles of a brain-computer interface (BCI) based on real-time functional magnetic resonance imaging (fMRI). *IEEE Transactions on Bio-Medical Engineering, 51*(6), 966-70.

Weiskopf, N., Scharnowski, F., Veit, R., Goebel, R., Birbaumer, N., & Mathiak, K. (2004). Self-regulation of local brain activity using real-time functional magnetic resonance imaging (fMRI). *Journal of Physiology, Paris, 98*(4-6), 357-73.

Weiskopf, N., Sitaram, R., Josephs, O., Veit, R., Scharnowski, F., Goebel, R., et al. (2007). Real-time functional magnetic resonance imaging: Methods and applications. *Magn Reson Imaging,*

Weiskopf, N., Veit, R., Erb, M., Mathiak, K., Grodd, W., Goebel, R., et al. (2003). Physiological self-regulation of regional brain activity using real-time functional magnetic resonance imaging (fMRI): Methodology and exemplary data. *NeuroImage, 19*(3), 577-86.

Wheeler, E. Z., & Fellows, L. K. (2008). The human ventromedial frontal lobe is critical for learning from negative feedback. *Brain, 131*(5), 1323.

Yucha, C., & Gilbert, C. (2004). *Evidence-based practice in biofeedback and neurofeedback*. Wheat Ridge, CO: Association for Applied Psychophysiology and Biofeedback. Donald Moss, Theodore J. LaVaque, and D. Corydon Hammond Series Editors.

Yoo, S. S., & Jolesz, F. A. (2002). Functional MRI for neurofeedback: Feasibility study on a hand motor task. *Neuroreport, 13*(11), 1377-81.

Yoo, S. S., O'Leary, H. M., Fairneny, T., Chen, N. K., Panych, L. P., Park, H., et al. (2006). Increasing cortical activity in auditory areas through neurofeedback functional magnetic resonance imaging. *Neuroreport, 17*(12), 1273-8.

APPENDICES

A. Funnelled Debriefing Interview

The following questions were verbally asked of subjects following scanning, in a semi-structured interview format. At times, the experimenter prompted subjects to elaborate. Subjects' responses were written down by the experimenter.

1) Tell me about your experience during the scanning.

2) Was the real-time feedback helpful or distracting? Did you prefer performing the task with or without the feedback?

3) Did your strategy change over the course of the experiment? If so, how? What strategies did you use?

4) Can you tell me what this experiment was about?

5) Did you notice a link between what you were thinking and what the feedback was showing?

6) Do you have any suspicions about the true purpose of this experiment?

B. Summary of Included and Excluded Subjects

Table B.1: Summary of included and excluded subjects

		Exclusion Criteria			
Subject Number	**Included**	Suspicion rating > 3	Motion > 3 mm	Task-motion corr. $r > 0.20$*	Brain abnormalities/ Poor normalization
S1	•				
S2		•			
S3	•				
S4		•			•
S5	•				
S6	•				
S7		•			
S8	•				
S9			•		
S10		•	•		
S11					•
S12	•				
S13	•				
S14				•	
S15	•				
S16	•				
S17		•			
S18	•				
S19	•				
S20	•				
S21	•				
S22		•		•	

* Correlation between PPA regulation task and motion parameter estimates

C. Individual Difference Measures

Table C.1: Mean scores for individual difference measures

	Current Sample		Previously Published	
	N = 13		(Benet-Martinez & John, 1998) *N* = 711	
BFI-44 Trait	**M**	**SD**	**M**	**SD**
Extraversion	3.45	0.46	3.2	0.8
Agreeableness	3.93	0.28	3.8	0.5
Conscientiousness	3.47	0.53	3.6	0.7
Neuroticism	2.08	0.72	3	0.8
Openness	4.04	0.44	3.7	0.6
	N = 13		(Trapnell & Campbell, 1999) *N* = 1137	
RRQ Scale	**M**	**SD**	**M**	**SD**
Rumination	3.46	0.65	3.46	0.71

D. Ethics Approval Certificate

The University of British Columbia
Office of Research Services
Clinical Research Ethics Board – Room 210, 828 West 10th Avenue, Vancouver, BC V5Z 1L8

ETHICS CERTIFICATE OF EXPEDITED APPROVAL: AMENDMENT

PRINCIPAL INVESTIGATOR:	DEPARTMENT:	UBC CREB NUMBER:
Kalina Christoff	UBC/Arts/Psychology, Department of	H06-03324

INSTITUTION(S) WHERE RESEARCH WILL BE CARRIED OUT:

Institution	Site
UBC	Vancouver (excludes UBC Hospital)

Other locations where the research will be conducted:
N/A

CO-INVESTIGATOR(S):

Irene Liu
Ronald Graeme McCaig

SPONSORING AGENCIES:

- Canadian Institutes of Health Research (CIHR) - "Investigating prefrontal cortex functions using real-time fMRI"

PROJECT TITLE:
Investigating prefrontal cortex functions using real-time fMRI

REMINDER: The current UBC CREB approval for this study expires: January 19, 2010

AMENDMENT(S):

AMENDMENT APPROVAL DATE:
April 14, 2009

Document Name	Version	Date
Protocol:		
Investigating real-time fMRI: Protocol with Revisions	2	March 26, 2009
Consent Forms:		
Consent form with Revisions - March 26	6	March 26, 2009
Other Documents:		
Deception form and verbal debriefing	1	March 26, 2009

CERTIFICATION:
In respect of clinical trials:
1. The membership of this Research Ethics Board complies with the membership requirements for Research Ethics Boards defined in Division 5 of the Food and Drug Regulations.
2. The Research Ethics Board carries out its functions in a manner consistent with Good Clinical Practices.
3. This Research Ethics Board has reviewed and approved the clinical trial protocol and informed consent form for the trial which is to be conducted by the qualified investigator named above at the specified clinical trial site. This approval and the views of this Research Ethics Board have been documented in writing.

The amendment(s) for the above-named project has been reviewed by the Chair of the University of British Columbia Clinical Research Ethics Board and the accompanying documentation was found to be acceptable on ethical grounds for research involving human subjects.

Approval of the Clinical Research Ethics Board by :

[redacted]

Dr. Caron Strahlendorf,
Associate Chair