Modulating reward learning in healthy adults with transcranial direct current stimulation: Towards a novel treatment for depression



Margot Juliëtte Overman

A thesis submitted in partial fulfilment of the requirements for the degree of Doctor of Philosophy at the University of Oxford.

University College Trinity Term 2018

“Through seeking and blundering we learn.”

Johann Wolfgang von Goethe

Modulating reward learning in healthy adults with transcranial direct current stimulation: Towards a novel treatment for depression

Margot Juliëtte Overman

Trinity Term 2018 DPhil in Clinical Neurosciences University College

Abstract

It has recently been proposed that deficits in reinforcement learning are a core feature of depressive disorders. This thesis aimed to investigate whether transcranial direct current stimulation (tDCS) could be applied to alter these learning processes in healthy adults and individuals with symptoms of depression. A first study assessed the psychometric properties of a computerised reinforcement learning task, which was the main paradigm used in the tDCS studies. Next, a neurostimulation study was designed to examine the effects of tDCS of the dorsolateral prefrontal cortex (DLPFC) on reinforcement learning parameters in healthy adults. It was found that prefrontal stimulation was associated with faster learning from positive but not negative outcomes when applied during task performance. The subsequent study investigated whether this effect of neurostimulation could be harnessed to augment a cognitive training paradigm aimed to enhance learning from either positive or negative events. No evidence was found that tDCS could boost effects of cognitive training on either reinforcement learning or affective biases in information processing in healthy adults. Finally, a pilot study examined whether tDCS of the DLPFC could affect reward learning in a sample of individuals with subclinical depression. While interpretations of the findings were limited by the study's small sample size, there were indications that active tDCS might enhance learning rates for both positive and negative outcomes during stimulation in this population. Altogether, this work presents preliminary evidence that non-invasive brain stimulation can transiently alter reinforcement learning processes in healthy adults and subclinical depression. Additional research is needed to examine the potential of this technique as a clinical intervention for depression and related psychiatric disorders.

Acknowledgements

My thesis is centred on the question of how people learn: what do we gain from our experiences, and how can we become better at what we do? The answer is perhaps most aptly given by Goethe's quote above, that "by seeking and blundering we learn" (an encouraging thought when doing a DPhil). Over the past few years, I have been given the opportunity to do a bit of blundering of my own, which has taught me more than I could have imagined when I first started my doctorate. I would hereby like to thank all the people who have supported me throughout this journey, and who have inspired me to never stop seeking.

First of all, a huge thank you to all the volunteers who were willing to sit through hours of listening to "cha-ching" and "buzz" sounds with drippy sponges strapped to their heads. This work would certainly not have been possible without your help, and your stories and insights kept me going through the many testing days.

For all the help I received during my DPhil I thank my supervisors Jacinta and Mike, who patiently answered my many questions (particularly when getting stuck with Matlab once again), and whose guidance and feedback were invaluable throughout the research and writing process.

I am thankful to my fellow FMRIB-ers for being there for all the piloting stages of my experiments, the weird lunchtime conversations, and generally making life in Oxford a lot more interesting. I've never seen a more collaborative or supportive place to work and feel very lucky to have been part of the Plasticity team.

To Owen, who has given me endless encouragement and support, especially in the final frantic stages of testing and writing. Thank you for putting up with my thesis mania, for bossing me into relaxing when I needed it, and most of all for being the best friend I could have wished for over the past two years. I can't wait to start our next adventure together.

Finally, I am incredibly grateful to my family, and especially my parents, who have been there for me all the way to and through this DPhil. To my dad, who has encouraged me to make my own choices, and who helped me to explore many foreign countries and finally find my place in England. And to my mom, for showing me how to be independent and adventurous, for the many phone calls during the hard times, and for her selfless support for as long as I can remember.

Table of Contents

Chapter 1. Non-invasive brain stimulation to change reward learning in depression: A computational approach....1

1.1. Introduction....1

1.2. A computational view of reinforcement learning....6

1.3. Reward learning deficits in depression....9
1.3.1. Reward and punishment sensitivity....10
1.3.2. Exploration and choice stochasticity....15
1.3.3. Adaptability and learning rates....19
1.3.4. Learning strategy and model selection....23
1.3.5. Overview of learning impairments in depression....24

1.4. Effects of NIBS on reinforcement-guided learning....25
1.4.1. Stimulation of the DLPFC alters reward sensitivity....26
1.4.2. NIBS can affect deliberate exploration....32
1.4.3. Excitation of the left DLPFC increases learning rates in healthy adults....35
1.4.4. Preliminary evidence for disruption of model-based learning with TMS of the DLPFC....38

1.5. Discussion....46
1.5.1. Overview of the findings....46
1.5.2. Future directions....47
1.5.3. Research aims and thesis outline....48

Chapter 2. Methods for stimulating neural networks of reinforcement learning....51

2.1. Introduction....51

2.2. Modulating the information content of events....51
2.2.1. Expected and unexpected uncertainty....51
2.2.2. The Information Bias Learning Task (IBLT)....52

2.3. Physiological mechanisms of Transcranial Direct Current Stimulation (tDCS)....53
2.3.1. General overview....53
2.3.2. Administration of tDCS....54
2.3.3. Physiological mechanisms of action....56
2.3.4. Spatial current distribution....61
2.3.5. Safety considerations....62

2.4. Rationale for the tDCS protocol....63
2.4.1. Neural target for stimulation....63
2.4.2. Stimulation parameters....65

Chapter 3. Examining test-retest reliability and practice effects on the Information Bias Learning Task (IBLT)....66

3.1. Introduction....66

3.2. Methods and Materials....68
3.2.1. Participants....68
3.2.2. General procedure....68
3.2.3. Information Bias Learning Task (IBLT)....69
3.2.4. Computational modelling and choice behaviour....72
3.2.5. Win-driven choice behaviour....76
3.2.6. Statistical analyses....77

3.3. Results 79
3.3.1. Task performance *79*
3.3.2. Test-retest reliability *80*
3.3.3. Practice effects *82*

3.4. Discussion 87
3.4.1. IBLT performance *87*
3.4.2. Test-retest reliability *87*
3.4.3. Practice effects *89*
3.4.4. Conclusion *91*

Chapter 4. Increasing learning rates for positive outcomes with bifrontal tDCS: A proof-of-concept study 92

4.1. Introduction 92

4.2. Methods and Materials 95
4.2.1. Participants *95*
4.2.2. General procedure *96*
4.2.3. Transcranial direct current stimulation (tDCS) *98*
4.2.4. Statistical analyses *99*

4.3. Results 100
4.3.1. Participant characteristics *100*
4.3.2. Study 1. Effects of online tDCS on reward learning in the IBLT *102*
4.3.3. Study 2. Effects of offline tDCS on reward learning on the IBLT *108*
4.3.4. Study 3. Replication of online tDCS effects *112*
4.3.5. Pooled data for the effects of online tDCS on learning rates *118*
4.3.6. Power calculations *121*
4.3.7. Participant blinding efficacy across studies *122*

4.4. Discussion 122
4.4.1. Main findings *122*
4.4.2. Study limitations *125*
4.4.3. Conclusion *127*

Chapter 5. Effects of concurrent tDCS and IBLT training on reward learning and affective biases in healthy adults 128

5.1. Introduction 128

5.2. Methods and Materials 133
5.2.1. Participants *133*
5.2.2. General procedure *133*
5.2.3. Facial Expression Recognition Task (FERT) *135*
5.2.4. Statistical analyses *136*

5.3. Results 137
5.3.1. Baseline participant characteristics *137*
5.3.2. Win-volatile IBLT training *138*
5.3.3. Loss-volatile IBLT training *142*
5.3.4. Comparison of Win-volatile and Loss-volatile training *148*
5.3.5. Participant blinding *153*

5.4. Discussion 153
5.4.1. Main findings *153*
5.4.2. IBLT training effects on reinforcement learning *154*
5.4.3. Impact of tDCS on learning parameters *156*
5.4.4. Affective biases in healthy adults *158*
5.4.5. Limitations and future directions *159*

5.4.6. Conclusion *160*

Chapter 6. Exploring the impact of tDCS over the DLPFC on reward-guided learning in subclinical depression **161**

6.1. Introduction 161

6.2. Methods and Materials 165
6.2.1. Participant recruitment *165*
6.2.2. Study design *166*

6.3. Results 166
6.3.1. Participant characteristics *166*
6.3.2. Effects of bifrontal tDCS on learning rates *167*
6.3.3. Effects of bifrontal tDCS on the inverse temperature *170*
6.3.4. Win-driven choice behaviour and total winnings *171*
6.3.5. Mood and anxiety measures *172*
6.3.6. Power calculation *173*
6.3.7. Participant blinding *173*

6.4. Discussion 174
6.4.1. Main findings *174*
6.4.2. Study limitations *175*
6.4.3. Conclusion *176*

Chapter 7. General discussion **177**

7.1. Summary of results 177
7.1.1. Study 1: Performance consistency on the Information Bias Learning Task *177*
7.1.2. Study 2: Prefrontal tDCS increases reward learning rates in healthy adults *180*
7.1.3. Study 3: Training on the IBLT is unaffected by tDCS in healthy adults *182*
7.1.4. Study 4: Global changes in learning parameters during tDCS in subclinical depression *184*
7.1.5. Main conclusions *185*

7.2. Methodological considerations 187
7.2.1. tDCS parameter optimisation *187*
7.2.2. Computational modelling of the tDCS current *188*
7.2.3. Double-blinding procedure *189*

7.3. Clinical relevance 191

7.4. Future directions 193

7.5. Final remarks 196

References **198**

Appendix A. Comparator models **244**

Appendix B. Mood and anxiety correlations **246**

Appendix C. Model-based and model-free analyses **249**

Appendix D. tDCS safety screening form **253**

Appendix E. Adverse effects questionnaire **254**

Appendix F. IBLT questionnaire **255**

Long abstract

Reinforcement learning is a fundamental adaptation mechanism, in which humans learn to optimise their choices based on prior experiences. Recently, it has been proposed that impairments in such learning processes may be a causal factor in the development of psychiatric disorders such as depression. Therefore, modulation of reinforcement learning mechanisms might present a novel pathway to improving symptoms of depressive disorders. In this thesis, it was investigated whether direct current stimulation (tDCS) of the dorsolateral prefrontal cortex (DLPFC) could change learning from positive and negative outcomes in healthy adults and individuals with subclinical depression.

Chapter 1 presents a review of the evidence for reinforcement learning impairments in depression from a computational psychiatry perspective. Next, research to date on the effects of non-invasive brain stimulation on different components of reward learning is described. Finally, it is proposed that regulating impairments in learning rates may be of particular interest in the treatment of depression.

Chapter 2 focuses on the theoretical principles of the Information Bias Learning Task (IBLT) and the physiology of transcranial direct current stimulation, the main cognitive task and neurostimulation technique applied in the present research. Furthermore, a rationale for the tDCS protocol applied in the present studies is provided.

Chapter 3 assesses test-retest reliability and practice effects on the IBLT, a computerised paradigm used to assess reinforcement learning in this thesis. It was found that reliability indices of learning rates on the IBLT ranged from low to moderate, depending on the specific block and outcome investigated. Intraclass correlation coefficients and standard error of measurements values are provided to facilitate interpretation of intervention studies using the IBLT. In addition, it was found that there

was a practice effect on the task, such that learning rates for negative outcomes were low in the first block but increased with repeated exposure to the task. The influence of potential practice effects should therefore be taken into account when interpreting changes of IBLT performance in repeated-measures studies.

Chapter 4 examines whether tDCS of the DLPFC at 2 mA for 20 minutes could change learning rates for positive and/or negative events on the IBLT in healthy adults. This proof-of-concept study demonstrated that stimulation applied during but not prior to task performance could increase the speed of learning for rewarding outcomes. In contrast, no changes in learning from punishments were observed with stimulation. This effect was not observed in a replication study, suggesting that the preliminary finding of an increase in reward learning with bifrontal tDCS must be interpreted with caution. However, pooled analyses demonstrated an increase in win learning rates with active stimulation across these two studies, providing some additional support for the hypothesis that bifrontal tDCS may be able to specifically change learning from positive information.

Chapter 5 investigates whether this effect of tDCS could be harnessed to enhance training towards learning from positive or negative events. Participants completed IBLT training towards learning from either win or loss outcomes while receiving sham or active tDCS of the dorsolateral prefrontal cortex. While there were minor indications that a negative training paradigm increased avoidance of negative outcomes, there were no significant effects of positive training on reinforcement learning. Furthermore, there was no evidence that tDCS could alter performance in either of the two training paradigms. It would be of interest to examine whether concurrent tDCS and IBLT training may have different effects in individuals presenting with baseline impairments in reinforcement learning.

Chapter 6 explores the effects of the tDCS protocol described in Chapter 4 on reward learning in a pilot study of individuals with subclinical depression. While interpretations from this study are curbed by the small sample size, preliminary results suggest that tDCS of the DLPFC may increase learning rates for both positive and negative outcomes during stimulation in this population.

Finally, Chapter 7 summarizes the findings reported in this thesis and considers future research directions based on these results. Overall, this work demonstrates that prefrontal tDCS may induce a transient increase in learning rates for positive outcomes. This technique could thus have the potential to modify hypothesised causal factors of mood disorders, and thereby provides an exciting early step in the development of novel treatments for depression.

This thesis contains approximately 48,700 words

List of key abbreviations

AIC	Akaike's Information Criterion
ANOVA	Analysis of variance
BART	Balloon Analogue Risk Task
BDI	Beck's Depression Inventory
BIC	Bayesian Information Criterion
cTBS	Continuous Theta Burst Stimulation
(d)ACC	(Dorsal) anterior cingulate cortex
DLPFC	Dorsolateral prefrontal cortex
FPC	Frontopolar cortex
FERT	Facial Expression Recognition Task
IBLT	Information Bias Learning Task
ICC	Intraclass correlation coefficient
MDD	Major Depressive Disorder
MEP	Motor evoked potential
NIBS	Non-invasive brain stimulation
OFC	Orbitofrontal cortex
PANAS	Positive and Negative Affect Schedule
(r)TMS	(Repetitive) Transcranial Magnetic Stimulation
SEM	Standard error of measurement
STAI	State-Trait Anxiety Index
tACS	Transcranial Alternating Current Stimulation
tDCS	Transcranial Direct Current Stimulation

List of Figures

Figure 1.1. The neurobiological effects of non-invasive brain stimulation.....5

Figure 1.2. Diagram of a simple reinforcement learning process.....9

Figure 1.3. Computational learning parameters associated with depression.....10

Figure 1.4. Neural foci of reward and punishment learning.....26

Figure 2.1. tDCS equipment and application.....55

Figure 2.2. Intra-stimulation effects of tDCS 58

Figure 2.3. Physiological after-effects of tDCS.....60

Figure 2.4. Electric field modelling at the gyri/sulci level 62

Figure 2.5. Connectivity and function of the dACC.....65

Figure 3.1. Schematic illustration of the Information Bias Learning Task 71

Figure 3.2. Order of the IBLT task blocks..... 71

Figure 3.3. Fit of comparator models 75

Figure 3.4. Schematic illustration of win-driven behaviour on the IBLT 77

Figure 3.5. Volatility effect on the IBLT.....80

Figure 3.6. Practice effects on the IBLT.....83

Figure 3.7. IBLT learning rates by run and session.....84

Figure 3.8. Inverse temperature values on the IBLT 85

Figure 3.9. Proportion of win-driven choices on the IBLT 86

Figure 4.1. Timeline of the tDCS procedure 98

Figure 4.2. Learning rates in Win- and Loss-volatile blocks with online tDCS.....104

Figure 4.3. Learning rates in Both-volatile blocks with online tDCS 106

Figure 4.4. Learning rates in Win- and Loss-volatile blocks with offline tDCS.....109

Figure 4.5. Learning rates in Both-volatile blocks with offline tDCS.....110

Figure 4.6. Learning rates in Win- and Loss-volatile blocks: Replication study 116

Figure 4.7. Learning rates in Both-volatile blocks: Replication study 117

Figure 4.8. Learning rates in Win- and Loss-volatile blocks for pooled data. 120

Figure 4.9. Learning rates in Both-volatile blocks for pooled data 121

Figure 5.1. Schematic representation of the IBLT training procedure. 135

Figure 5.2. Learning rates in Win-volatile training blocks. 140

Figure 5.3. Facial Emotion Recognition Task after Win-volatile training. 141

Figure 5.4. Learning rates in Loss-volatile training blocks. 143

Figure 5.5. Inverse temperature values in Loss-volatile training blocks 144

Figure 5.6. Win-driven choices in Both-volatile blocks with Loss-volatile training. 146

Figure 5.7. Facial Emotion Recognition Task after Loss-volatile IBLT training 147

Figure 5.8. Learning rates in Win- and Loss-volatile training.....149

Figure 5.9. Win-driven choices in Both-volatile blocks with Loss-volatile training 152

Figure 5.10. Win-driven choices in Both-volatile blocks by training type.....152

Figure 6.1. Learning rates in Win- and Loss-volatile blocks.....169

Figure 6.2. Inverse temperature values in Win-and Loss-volatile blocks 170

Figure 6.3. Win-driven choices in Win- and Loss-volatile blocks 172

Figure 6.4. PANAS Positive scores with tDCS 173

Figure 7.1. Simulation of tDCS current 189

Figure 7.2. Illustration of 1 mA electric field strength 189

List of Tables

Table 1.1. tDCS and TMS studies on reinforcement learning components 41

Table 3.1. Description of comparator models of the IBLT. 75

Table 3.2. Test-retest reliability estimates of win and loss learning rates 81

Table 3.3. Test-retest reliability estimates of win – loss learning rates.....81

Table 3.4. Test-retest reliability estimates of volatile – stable learning rates.....81

Table 3.5. Test-retest reliability estimates of the inverse temperature 82

Table 3.6. Test-retest reliability estimates of win-driven choice behaviour.....82

Table 4.1. Baseline anxiety and mood scores with online/offline tDCS 102

Table 4.2. Baseline anxiety and mood scores for online/replication study 114

Table 5.1. Baseline anxiety and mood scores by training group 138

Table 5.2. Linear mixed model of computational parameter values. 146

Chapter 1. Non-invasive brain stimulation to change reward learning in depression: A computational approach

1.1. Introduction

Decision making is a complex process, which requires integration of currently available sources of information with knowledge gained from previous experiences. For example, when choosing whether to travel to France or Scotland when going on holiday, you might consider the most recent weather predictions and costs of the two different destinations, as well as take into account how much you enjoyed spending time in the respective countries in the past. In humans and other animals, choices are often driven by beliefs about how rewarding the outcome of a particular action is likely to be. If a choice leads to a more positive or negative result than expected (for example, if vacationing in France happens to be more expensive than predicted), estimates of the value of that action tend to change accordingly. In formal terms, this process of updating action-outcome associations is known as *reinforcement learning*, in which individuals can maximise gains and minimise punishments by learning the effects of different behaviours in specific environments (Sutton & Barto, 1998).

Recently, it has been proposed that deficits in reinforcement learning and decision making are a core feature of psychiatric disorders such as Major Depressive Disorder (MDD) (Huys, Daw, & Dayan, 2015; Moutoussis, Story, & Dolan, 2015; Whitton, Treadway, & Pizzagalli, 2015). In general, individuals with depression are thought to be hyposensitive to rewards, while having overreactive responses to punishments and/or having difficulties in utilising negative feedback to improve performance (Eshel & Roiser, 2010). As a result, there may be an imbalance in how much depressed individuals learn from positive and negative information. Impairments in reinforcement learning can have serious consequences, as an inability to modulate behaviour according to prior experiences

may lead to suboptimal choices and therefore a greater chance of encountering adverse outcomes in the future. This, in turn, could strengthen negative beliefs held by depressed patients and thereby induce a vicious circle. In this thesis, I therefore propose that it may be possible to improve symptoms of depression by altering reinforcement learning processes.

A potential method for ameliorating depressive symptomatology, by changing causal cognitive and behavioural processes, is the use of non-invasive brain stimulation (NIBS) techniques. A main aim of these methods is to alter specific patterns of neural activity associated with psychiatric disorders in a manner which produces symptom reduction. The majority of NIBS studies targeting depression have focused on changing activity in the left dorsolateral prefrontal cortex (DLPFC). There is evidence from neuroimaging studies that the left DLPFC is hypoactive in depression (Disner et al., 2011; Siegle et al., 2007), which has been linked to negative emotional judgement (Grimm et al., 2008) and dysfunctional regulation of negative affect (Koenigs & Grafman, 2009). NIBS studies have therefore emphasised the potential of reducing depressive symptoms by increasing activity within this region. The most commonly used forms of NIBS for depression are Transcranial Magnetic Stimulation (TMS) and Transcranial Direct Current Stimulation (tDCS). In TMS, a rapidly changing magnetic field is applied to the scalp to induce electrical currents in targeted brain regions (see Figure 1.1.A.; Priori, Hallett, & Rothwell, 2009). The effect of the stimulation on neural excitability is dependent on a variety of factors, including the frequency, timing, and intensity of the magnetic pulses (Hallett, 2000; Sandrini, Umiltà, & Rusconi, 2011). In addition to modulating cortical excitability during stimulation, it has been shown that repetitive TMS (rTMS) protocols can be used to increase spontaneous neural firing outlasting the period of stimulation (see Figure 1.1.B.). When applied over the left DLPFC daily for 4 to 6 weeks at high

frequencies, rTMS has consistently been found to reduce symptoms of depression (Berlim et al., 2014; Gaynes et al., 2014; Perera et al., 2016). The clinical efficacy of TMS was verified in a large multisite trial (O'Reardon et al., 2007), which resulted in approval from the U.S. Food and Drug Administration (FDA) of this protocol as a therapy for treatment-resistant depression in 2008.

Similarly, there has been accumulating evidence suggesting positive effects of tDCS on mood disorders, although the effectiveness of this approach is less well established than in TMS. In tDCS, a weak electric current of 1-2 mA is delivered through electrodes placed on the scalp in sponges (see Figure 1.1.C.; Priori, Hallett, & Rothwell, 2009). Whereas TMS can directly affect neuronal firing rates, tDCS is thought to change cortical excitability by altering resting membrane potentials during stimulation, as well as inducing after-effects outlasting the stimulation through synaptic modification (Nitsche et al., 2003; Nitsche & Paulus, 2000; see Figure 1.1.D. and 1.1.E.). Anodal stimulation increases neural excitation, whereas cathodal stimulation inhibits neural activation (Figure 1.1.F.; Nitsche & Paulus, 2000; see Chapter 2 for a more extensive discussion on the physiological basis of tDCS effects). Studies of depression tend to employ a bifrontal tDCS montage, in which the anodal electrode is positioned over the left DLPFC and the cathodal electrode is placed over the right DLPFC. This choice of the cathodal electrode position is mainly driven by research indicating that depressed individuals show a hemispheric imbalance of DLPFC activation, such that the left DLPFC is hypoactive and the right DLPFC is hyperactive (Disner et al., 2011; Grimm et al., 2008; Hecht, 2010). Randomised controlled studies of bifrontal tDCS of the DLPFC have reported greater reductions in depression with active compared to sham tDCS (Brunoni et al., 2011a; Dell'Osso et al., 2012; Ferrucci et al., 2009) or occipital tDCS (Boggio et al., 2008). In addition, a meta-analysis of individual patient data from 289 participants showed that

active bifrontal tDCS was more effective in improving depressive symptoms than sham tDCS, with remission and response rates being similar to those found with TMS (Brunoni et al., 2016). However, it should be noted that not all research to date has supported strong antidepressant effects of prefrontal tDCS, with several placebo-controlled trials failing to find greater remission with active stimulation compared to sham stimulation in treatment-resistant depression (Blumberger et al., 2012; Palm et al., 2012). Furthermore, it has been suggested that the effect size of the improvements in depression scores is small to medium (Shiozowa et al., 2014). Additional research is therefore needed to assess and further improve the clinical efficacy of bifrontal tDCS.

In addition to moderating depressive symptoms directly, NIBS techniques present an opportunity for altering specific behaviours and cognitions which are causally related to depression by changing activity in key neural substrates underlying these functions. Previous reviews of NIBS effects have indicated that tDCS and TMS may be capable of influencing a wide range of processes related to decision making, such as risky behaviour and impulsivity (Brevet-Aeby et al., 2016; Levasseur-Moreau & Fecteau, 2012). In addition, NIBS has been shown to transiently affect general cognitive functions such as working memory (Brunoni & Vanderhasselt, 2014) and attention (Roy et al., 2015), which play a significant role in complex decision making. However, it is at present unclear whether TMS and tDCS can be applied to modulate reinforcement learning processes related to psychiatric disorders. In this literature review, I aim to explore whether NIBS is an effective method for altering reward learning mechanisms relevant to depression. First, I will provide a brief outline of common mathematical models used in research of reinforcement learning. Next, I will discuss behavioural and neural evidence implicating learning impairments in the development and maintenance of depressive disorders. It has previously been suggested that applying a computational framework to such deficits can

provide a means of disentangling multiple processes which may be involved concurrently in decision making (Scholl & Klein-Flügge, in press). I will therefore particularly focus on evidence describing learning impairments from a computational psychiatry perspective. In the remainder of the review, I will examine research studies on the use of NIBS in modulating reward learning in both healthy adults and individuals with symptoms of depression. Finally, I will summarise the main aims of the present thesis informed by the findings from this literature review.

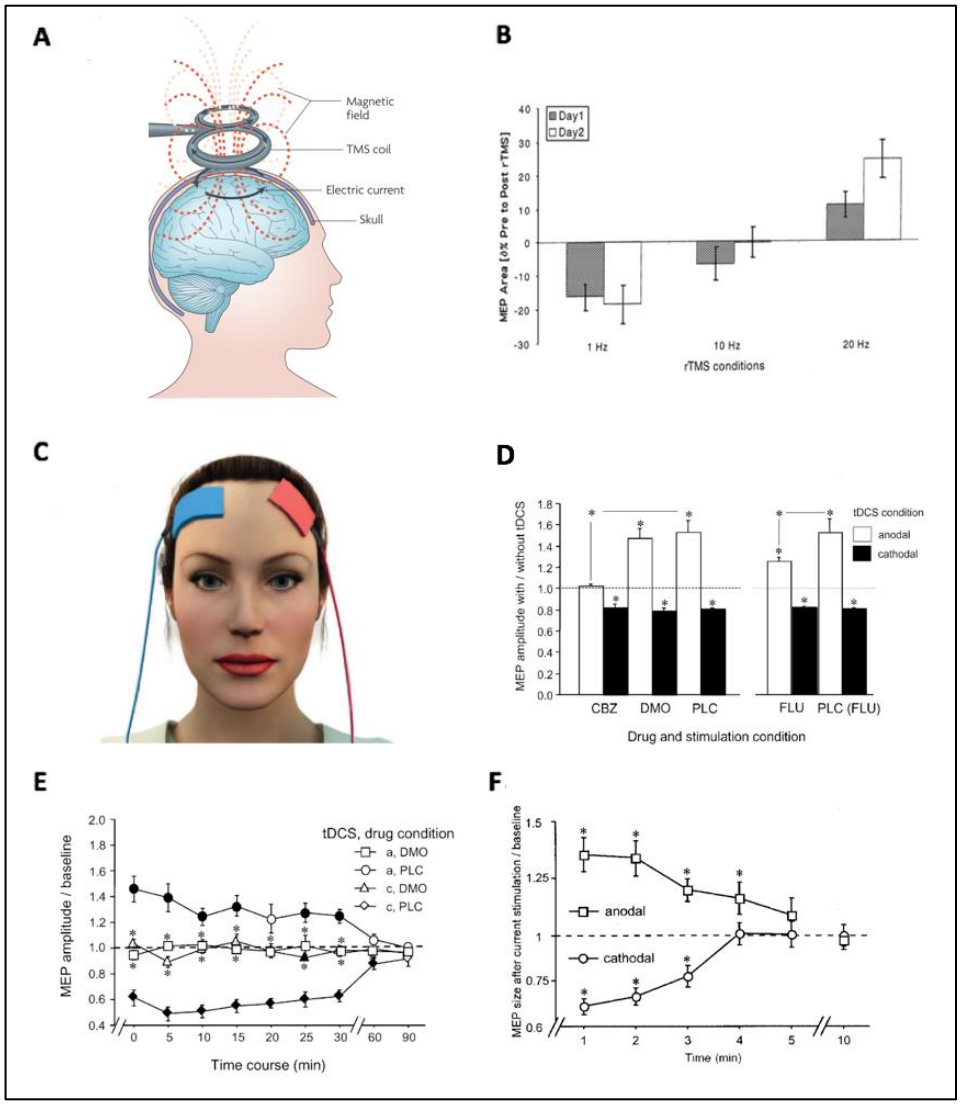

Figure 1.1. The neurobiological effects of Transcranial Magnetic Stimulation (TMS) and Transcranial Direct Current Stimulation (tDCS). **(A)** Illustration of the electrical currents produced by pulses of TMS, reproduced from Ridding and Rothwell (2007). **(B)** Percentage change in motor

evoked potentials (MEP) from pre- to post-rTMS application with 1 Hz, 10 Hz, and 20 Hz TMS to the motor cortex on two successive days. The findings indicate that low-frequency TMS inhibited corticospinal excitability, whereas stimulation at 20 Hz increased excitability (from Maeda et al., 2000). **(C)** Illustration of a tDCS montage targeting the left dorsolateral prefrontal cortex (anode; red) and the right dorsolateral prefrontal cortex (cathode; blue) reproduced from Filmer et al. (2014). The current flows from the anode to the cathode, thereby changing polarization of the underlying cortical tissue. **(D)** Graph indicating the proposed mechanisms for the immediate effects of tDCS (from Nitsche et al., 2003). The increase in cortical excitability resulting from anodal tDCS is abolished with administration of sodium channel blocker carbamezipine (CBZ) and diminished after supplementation of calcium channel blocker flunarizine (FLU), whereas NMDA receptor antagonist dextromethorphane (DMO) and placebo drugs (PLC) had no effects on excitability during tDCS. This suggests that intra-stimulation effects of tDCS are due to alterations of the resting membrane potential rather than changes in synaptic plasticity. **(E)** Graph demonstrating the mechanisms underlying after-effects of tDCS (from Nitsche et al., 2003). With both anodal (a) and cathodal (c) tDCS, excitability changes following stimulation are eliminated with administration of DMO but not placebo drugs (PLC), indicating a crucial role for NMDA-receptor-dependent synaptic modulation in the generation of after-effects of tDCS. **(F)** Polarity-specific effects of tDCS on cortical excitability in the motor cortex as described by Nitsche and Paulus (2000). The size of the motor-evoked potential (MEP) represents the excitability of the motor system. With 5 minutes of anodal tDCS at 1 mA over the motor cortex excitability was increased, whereas the opposite effects were observed with cathodal tDCS. Amplitudes returned to baseline within 5 minutes after cessation of the stimulation.

1.2. A computational view of reinforcement learning

The field of computational psychiatry is heavily influenced by the work of David Marr (1982) on visual perception, which proposes that information processing systems such as the brain can be understood at three complementary levels of analysis. The *computational* level specifies the task or problem that the system is designed to solve. The *algorithmic* level specifies how the system can solve the problem in mathematical terms. Finally, the

implementational level specifies how such processes can be realised physically. It has been argued that this approach can also be applied to clarify the mechanisms underlying processes central to psychiatric disorders, such as mood and belief formation (Kendler, 2008). There are several ways in which computational modelling can contribute to our understanding of cognitive and behavioural impairments in depression. Firstly, models can provide normative values against which impaired behaviour can be tested and quantified (Chen et al., 2015). In addition, they can enhance our ability to objectively measure different types of behaviour and formulate testable quantitative hypotheses (Adams, Huys, & Roiser, 2016). Finally, as previously discussed, formal models can test the individual contribution of different processes during the performance of multifaceted tasks (Scholl & Klein-Flügge, 2017), thereby providing a richer understanding of which specific parameters are most strongly involved in producing cognitive deficits in depression.

A central tenet of the present thesis is that depressive disorders may be characterized by an impairment in reinforcement learning. At a computational level, reinforcement learning can be described as the task of improving behaviour to maximise total or average rewards, while minimising the chance of receiving negative outcomes (Doya, 2007). In order to make an optimal choice, people have to consider which of the available information is relevant, evaluate their options, and decide which action should be taken (see Figure 1.2.). At the algorithmic level, a variety of models have been developed to study the way in which people can update their understanding of action-outcome associations. The most commonly used model, which is known as the Rescorla-Wagner model (Rescorla & Wagner, 1972), suggests that learning is not driven by absolute outcomes *per se*, but is rather the result of differences between the actual outcome and the predicted outcome, which is mathematically represented as a *prediction error* ε. The impact of the prediction error on current estimations of the *associative strength* ($r_{(i)}$) of

the action and outcome is then weighted by a *learning rate* α, which can differ across individuals, environments, and stimuli (Rescorla & Wagner, 1972; Sutton & Barto, 1998):

$$r_{(i+1)} = r_{(i)} + \alpha\varepsilon_{(i)} \quad (1)$$

with $r_{(i+1)}$ representing the updated associative strength (range 0 to 1).

In addition to specifying how action or stimuli values are updated, many learning models incorporate a function which describes how the results of internal evaluations are converted to actions (Huys, Maia, & Frank, 2016). A frequently applied function is the softmax, which utilises a Boltzmann distribution to favour actions with higher estimated values (Doya, 2002; Ishii, Yoshida, & Yoshimoto, 2002; Tokic, 2010; Whiteson, Taylor, & Stone, 2007):

$$Action\ Probability\ \{action \mid state\} = \left[\frac{e^{\beta Q(\text{state,action})}}{\sum_b e^{\beta Q(\text{state,action}')}}\right] \quad (2)$$

where Q(state, action) denotes the expected reward when selecting the action in a specified state, and β signifies a parameter named the inverse temperature. A value of 0 for the inverse temperature means that all options are chosen with equal probability independent of their value, while a large inverse temperature suggests that choices closely follow the estimated value of the various available options (Doya, 2008). In other words, individuals with a high inverse temperature value tend to choose the option that they believe is most likely to result in the best possible outcome. By applying these computational calculations in studies of depression, we can gain greater insight into how parameters and their associated behaviours are implemented and impaired in MDD.

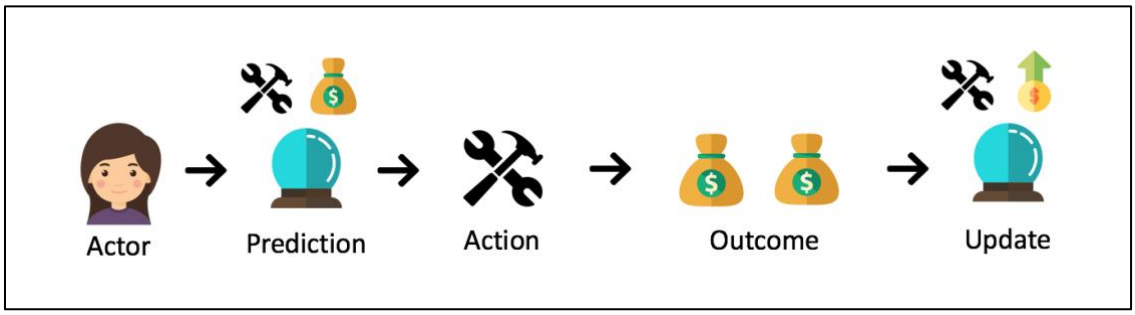

Figure 1.2. Diagram representing the core components of a simple reinforcement learning process. The actor decides to take a specific action based on a predicted outcome. If the outcome is better than expected, the estimated value of the action is increased, thereby increasing the likelihood that the action will be repeated in the future. Conversely, if the outcome is worse than expected, the estimated value of the action will be reduced and the likelihood of action repetition decreases.

1.3. Reward learning deficits in depression

There is a substantial evidence base implying a presence of reward-based learning deficits in depressive disorders. For instance, participants with depression have been found to choose highly-rewarded options in a probabilistic learning task less frequently than healthy controls (Kunisato et al., 2013). Similarly, healthy individuals tend to acquire a response bias towards highly rewarded stimuli, whereas MDD patients fail to develop this bias over time (Pizzagalli et al., 2008). It is possible that impairments of learning in depression depend on the valence of the task stimuli used. For example, the literature suggests that, in probabilistic learning tasks, depressed people are better at minimising punishments than healthy controls, whereas they tend to underperform in tasks where the aim is to maximise rewards (Beevers et al., 2013; Maddox et al., 2012). It can therefore be hypothesised that MDD patients are unable to use reward outcomes sufficiently to inform their actions, whereas they can present with increased responses to punishments or losses compared with non-depressed individuals. When looking at reinforcement learning through a

computational lens, such alterations in behaviour can be understood as resulting from deficits arising at various stages of the learning process (see Figure 1.3).

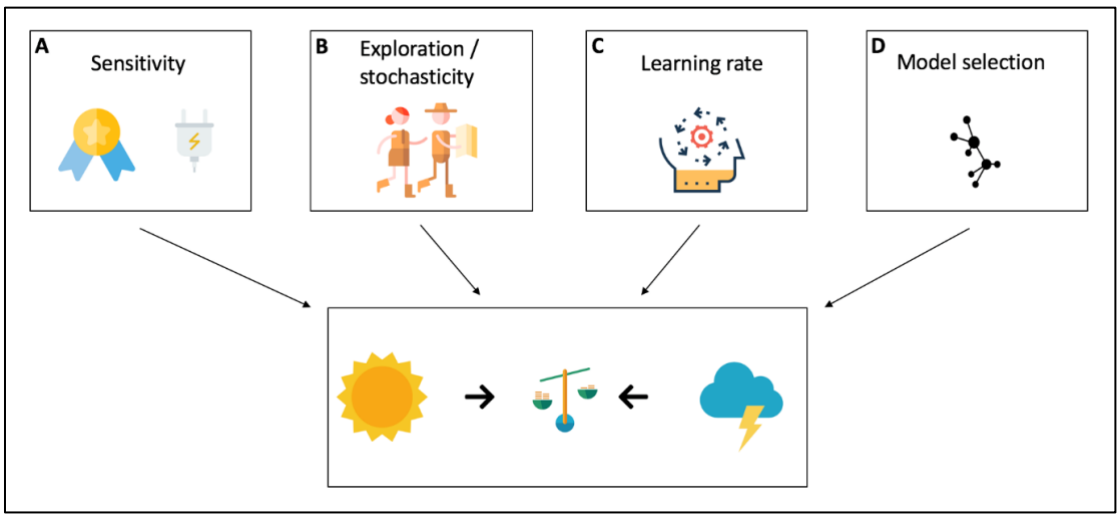

Figure 1.3. Potential parameters contributing to suboptimal learning from positive and negative events in depression. **(A)** Reduced sensitivity to positive outcomes and/or heightened sensitivity to aversive outcomes may result in an imbalance in learning from positive and negative feedback. **(B)** A failure to utilise the estimated values to inform actions can result in inappropriate exploratory or stochastic behaviour, thereby leading to suboptimal choices. **(C)** How quickly estimated values are updated based on experience is determined by the learning rate. If individuals cannot adapt their expectations correctly based on the information received, learning becomes impaired. **(D)** An inability to construct complex models of the world (model-based learning) in favour of simplistic models (model-free learning) may result in an inability to gain an understanding of intricate action-outcome associations.

1.3.1. Reward and punishment sensitivity

The Diagnostic and Statistical Manual of Mental Disorders 5 (DSM-5) indicates that a main characteristic of depressive disorders is a decrease in the enjoyment of pleasurable activities or anhedonia (American Psychiatric Association, 2013). A first explanation of impaired reward learning may therefore be that depressed individuals value rewards to a lesser extent or, alternatively, are more sensitive to adverse outcomes than healthy adults. Sensitivity refers to the degree to which behaviour is motivated by rewards or punishments

(Gray & McNaughton, 2000). An imbalance in sensitivity to positive and negative outcomes can lead to a diminished capacity to pursue and react to rewards (Alloy et al., 2016), which could consequently manifest in biased learning mechanisms and various depressive symptoms. For example, it has been proposed that reduced sensitivity to rewards may lead to decreases in motivation, a commonly observed symptom associated with depression (Treadway et al., 2012). Enhanced sensitivity to punishments, on the other hand, might play a role in producing avoidance and withdrawal behaviours associated with depression (Trew, 2011).

In terms of primary rewards, there is limited evidence that the basic sensory experience of rewarding stimuli is altered in depression. For instance, subjective ratings of pleasure in response to sweet taste do not differ between healthy adults and depressed individuals (Dichter et al., 2010; Thomsen, 2015). Reduced responsiveness has been observed for primary adverse outcomes, on the other hand, with depressed individuals having higher thresholds for painful stimuli than non-depressed controls (Dickens et al., 2003). Given the strong indications that depressed patients are biased towards processing negative information (e.g. Peckham et al., 2010), this decrease in sensitivity to aversive stimuli is counterintuitive. However, blunted responses to negative outcomes might explain why some studies have shown that depression was associated with reduced adaptation of behaviour following adverse outcomes (Eshel & Roiser, 2010). While findings from behavioural studies thus do not follow the expected pattern of hyposensitivity to rewards and/or hypersensitivity to punishments, neuroimaging studies have demonstrated that differences in neural patterns for positive and negative outcomes may differ between healthy controls and individuals at risk for depression in the absence of behavioural indications. Young people with a family history of depression were found to have blunted responses in the orbitofrontal cortex (OFC) following the presentation of

pleasant tastes or images compared to healthy participants, while increased activation was observed in the lateral OFC and insula in response to aversive stimuli (McCabe et al., 2012). In addition, recovered MDD patients have been found to have reduced blood-oxygen-level dependent (BOLD) responses in the ventral striatum when presented with primary rewarding stimuli, whereas increased activity was found in the bilateral caudate compared with controls when viewing unappetising images (McCabe et al., 2009). These findings suggest that, while subjective changes in the experience of pleasant or unpleasant stimuli may be lacking, altered neural responses could reflect an endophenotype for the risk of developing a mood disorder.

While there are therefore some indications that primary reward and punishment sensitivity is altered in depression, these studies have not investigated how such alterations might affect reinforcement learning. Studies using secondary rewards and punishments such as monetary reimbursement, however, have provided more insight into potential associations between sensitivity and reward learning deficits. Such behavioural changes have most frequently been investigated with the use of risky choice paradigms. It has been suggested that, in risky environments, MDD patients may be exceptionally sensitive to increases in losses while being hyposensitive to changes in rewards. This could then result in depressed participants taking riskier choices in order to avoid punishments than healthy controls, while becoming less likely to make risky decisions when trying to obtain rewards (Chen et al., 2015). However, findings of changes in risky behaviour in both reward- and loss-contexts have been mixed in depressed populations. On the commonly employed Iowa Gambling Task (Bechara et al., 2007), participants have to choose cards from four different decks, two of which are associated with disadvantageous monetary outcomes in the long term. In one version of the task (Version A), the disadvantageous decks have high immediate rewards but even higher punishments. In a second version of the task (Version

B), disadvantageous decks contain small punishments but even smaller rewards. Must and colleagues (2006) found that participants with depression chose cards from disadvantageous decks more frequently than healthy participants when presented with task Version A. Similarly, a study utilising a contingency-shifting variant of Version A, in which associations between stimuli and outcomes progressively shift during three phases of the task, showed that patients with MDD made poorer choices on both the regular task and the contingency-shift phases compared to healthy adults (Cella, Dymond, & Cooper, 2010). When using task Version B, on the other hand, no differences were found between depressed and non-depressed individuals in task performance (Must et al., 2006). Contrary to expectations, these findings suggest that depressed patients may be particularly sensitive to rewards, while responses to immediate punishments are unchanged. In contrast, however, Smoski et al. (2008) reported that participants with depression tended to choose fewer risky cards on the Version A task, indicating greater sensitivity to adverse outcomes. In addition, in a task where participants were required to exert a certain amount of effort in order to gain a reward, it was found that patients with MDD altered their choices to a lesser degree in response to changes in reward magnitude than healthy controls (Treadway et al., 2012), pointing to a reduction in sensitivity to positive events.

While there is thus some evidence to support the presence of hypersensitivity to punishments and hyposensitivity to rewards in association learning, the findings overall are rather mixed. It should be noted, however, that a non-trivial complication of tasks such as the Iowa Gambling Task is that differentiating between reward sensitivity and closely related constructs can be challenging. For example, impaired performance on this task could also be explained by reduced learning rates from the information provided, or by increases in exploratory behaviour leading to suboptimal choices. To date, only one study has explicitly compared the influence of reward sensitivity and learning rates on reward

learning behaviour in depression. In this study, data from 392 sessions of a Probabilistic Reward Task were re-analysed using computational reinforcement learning models (Huys et al., 2013). In the Probabilistic Reward Task, participants are presented with "rich", frequently rewarded stimuli as well as "lean", rarely rewarded stimuli. It has previously been found that healthy participants tend to develop a response bias to the rich stimuli over time on this task, whereas the response bias is absent in depressed individuals (Pizzagalli, Jahn, & O'Shea, 2005). In the meta-analysis, it was shown that this lack of a response bias was mainly related to reduced sensitivity to rewards in MDD and anhedonia, whereas learning rates were unaffected during task performance (Huys et al., 2013). As this task does not involve punishments, however, responses to negative outcomes could not be investigated.

Similar to the findings for neural responses to primary reinforcers, neuroimaging studies involving secondary rewards and punishments broadly support the hypothesis that reward sensitivity is altered in depression. In MDD patients, reduced reward prediction error signals have been observed in the striatum, medial orbitofrontal cortex, and midbrain compared to controls (Gradin et al., 2011; Rothkirch et al., 2017). While these alterations might reflect difficulties in accurately estimating the difference between reward predictions and outcomes, they could also represent general blunting of responses to rewards. Findings for changes in punishment-based conditions appear to be relatively more complex. While some studies have reported no differences in the neural encoding of punishment prediction errors (Rothkirch et al., 2017), others have demonstrated a hypersensitivity to negative outcomes in depression. When asked to choose between two options which were probabilistically associated with adverse outcomes, for example, greater neural signals in mediofrontal regions were observed following decision making errors in MDD patients compared to healthy controls (Cavanagh et al., 2011). Moreover,

these differences were associated with increased avoidance learning in depression. Of relevance, the differences in mediofrontal activity only occurred in trials in which participants had to choose between two stimuli which had nearly equal probabilities (e.g., 20% versus 30%) of resulting in a loss. In trials in which the difference in probabilities was large, on the other hand (e.g., 20% versus 80%), which was similar to the task paradigm employed by Rothkirch et al. (2017), no differences in neural activity were observed. It is therefore possible that the neural mechanistic changes associated with punishment sensitivity in depression are subtle and dependent on task demands.

While the findings to date have been heterogeneous, there is preliminary evidence to suggest that depression might be related to blunted sensitivity to rewarding stimuli and hypersensitivity to adverse outcomes. Disparities in findings across studies are likely due in part to task differences, such as the use of primary versus secondary reinforcers, provision of simple sensory stimulation instead of the application of reinforcement learning tasks, paradigms focusing on either rewards or losses in isolation, and the contribution of traits unrelated to the construct of interest to depression (e.g. impulsivity) to task performance. These aspects of the task design therefore need to be considered carefully when investigating alterations in sensitivity in depressed populations. In addition, many tasks used in the research discussed here are unable to distinguish reward sensitivity from alternative learning parameters, leaving the question open whether these are responsible for observed changes in task performance. Further investigations are therefore needed to clearly establish deviations in sensitivity from processes such as exploration and speed of learning.

1.3.2. Exploration and choice stochasticity

An alternative explanation of reinforcement learning deficits may be that depressed patients utilise their knowledge of the expected value of possible actions to a different

extent than healthy individuals when making a decision. In a simplistic world, the most straightforward strategy for optimising rewards is to consistently opt for the action which is thought to have the greatest chance of resulting in a rewarding outcome. In real life, however, people tend to respond in a much more variable manner, such that options which are estimated to have a lower likelihood of resulting in a positive outcome are occasionally chosen as well. This variability can either be the product of a deliberate decision to explore the available options, or from a failure to apply the learned value to choice behaviour, thus leading to an increase in random choices (Doya, 2008). In computational models, this behaviour is captured by the inverse temperature parameter as described in Equation 2.

Previously, it has been proposed that patients with depression might display exploratory or random behaviour to a different extent than healthy individuals. For example, it has been suggested that pessimistic people may engage less in exploratory behaviour than optimistic individuals, as they have more negative predictions of the value of the unexplored options (Huys, Daw, & Dayan, 2015). Surprisingly, however, some research has indicated that depressed individuals may present with increased rather than decreased exploratory behaviour. For instance, increased exploration has been reported in depressed individuals when performing the Leapfrog Task, in which one of two options results in a higher reward than the other. At any time, the reward outcome for the suboptimal choice can increase, thereby turning into the better option. In this task, consistently choosing the same option is a maladaptive strategy, as it can lead to changes in reward associations going undetected. Continuously switching between the two options, on the other hand, means that the best choice is not exploited sufficiently and can similarly result in lower overall winnings. The latter has been found to occur in depressed participants, who failed to exploit the optimally rewarding option when completing the

Leapfrog Task (Blanco et al., 2013). This finding contradicts the suggestion that depressed individuals may be averse to exploring novel options.

It should be noted, however, that exploratory or random behaviour can in practice also be influenced by sensitivity to rewards and punishments (Huys et al., 2013). It is particularly difficult to distinguish between these two parameters in commonly used two-alternative forced choice tasks, in which participants are required to choose one of two options. If the participant is less sensitive to reward outcomes and therefore values a less-rewarded option to the same extent as more frequently rewarded options, this could result in random responding on the task. Alternatively, the participant may have higher estimated values for one option, but fail to utilise this value correctly, similarly resulting in random responses or seemingly exploratory behaviour (Robinson & Chase, 2017). It can be hypothesised, however, that exploration caused by altered sensitivity should depend on the valence of the presented stimuli. As described in the previous example, individuals who are hyposensitive to positive outcomes might use rewards to a smaller extent in guiding their choices and thus show more exploratory behaviour. In contrast, hypersensitivity to negative outcomes might result in increased punishment-guided decision making, such that choices closely follow observations of adverse events. Of interest, in tasks where the aim is to minimise punishments rather than maximise gains, on the other hand, depressed participants have been found to switch between actions less frequently than non-depressed controls, which resulted in greater exploitation of the superior choice (Beevers et al., 2013). While this behaviour might be advantageous for enhancing task performance within the constraints of experimental paradigms, decreased exploration following negative outcomes in realistic situations could curtail opportunities to discover changes in the environment, or prevent learning that an adverse outcome might have been a chance event. This could consequently lead to perseverance of maladaptive behaviours and cognitions.

Combined with the findings of reduced exploitation of rewarding outcomes in the Leapfrog Task (Blanco et al., 2013), these results of increased exploration in punishment conditions suggest that changes in exploratory behaviour in depression might be due to hypersensitivity to punishments and hyposensitivity to rewards as outlined in Section 1.3.1. However, the analyses employed in these studies do not allow for a clear distinction between the two factors. Some studies have attempted to investigate changes in exploration more formally by taking a computational modelling approach. Overall, this research has also provided little evidence that depressive symptoms are related to differences in exploration or stochastic behaviour. For instance, a functional magnetic resonance imaging (fMRI) study found that, when thirsty participants were asked to choose between two stimuli in order to maximise the delivery of water drops, there were no differences in inverse temperature values between healthy adults and depressed individuals (Gradin et al., 2011). Similarly, in a probabilistic learning task involving both positive and negative stimuli, depressed individuals showed similar exploration/exploitation parameter values compared with healthy controls when acquiring knowledge of the associations between actions and outcomes (Chase et al., 2010). In light of these findings, it can be hypothesised that exploratory behaviour or choice stochasticity may be mainly driven by alterations in sensitivity, leading to increased exploration in reward contexts and decreased exploration in punishment conditions in depressed individuals compared with healthy adults. Importantly, however, studies to date have not attempted to explicitly differentiate exploratory behaviour from hyper- or hyposensitivity to negative and positive outcomes, respectively. In future studies, this might be accomplished by the development of carefully designed experimental tasks and/or through the application of computational models incorporating both sensitivity and exploration parameters.

A further gap in the literature currently available is that research studies have not distinguished between deliberate exploration and randomness of choice, which could be differentially affected in depression. For instance, specific impairments in deliberate exploration may be expected in patients with depression as a result of overly pessimistic predictions of the outcomes of alternative choices. Such changes in deliberate exploration might thus reflect a misconstruction of the state of the environment. Changes in randomness of choices, on the other hand, could point to a more basic deficit in utilising the estimated values. Additional research exploring potential differences in these two distinct processes may provide greater insight into the underlying cognitive mechanisms leading to depressive symptoms. Non-invasive brain stimulation studies aiming to induce changes in exploratory behaviour should therefore be aware of the potential influence of reward sensitivity as well as distinctions between deliberate and random exploration, and ideally utilise tasks which can differentiate between these individual parameters.

1.3.3. Adaptability and learning rates

A third possibility is that depressed individuals are sensitive to and act appropriately on rewards and punishments, but fail to update the value associated with different actions correctly based on prior experiences. In order to adjust estimated values accurately in response to an outcome, an actor needs to determine how informative the event is. If the event is thought to provide a lot of information regarding the action-outcome association at hand, estimates should be changed to a great degree. Conversely, if the event carries little information, estimates should be updated a small amount if at all. In reinforcement learning frameworks, the impact of single events on people's value estimates is reflected in the product of their learning rates and prediction errors (see Equation 1), with high learning rates indicating a belief that the observed prediction error is highly informative.

Previously, it has been shown that healthy individuals can maintain separate

learning rates for positive and negative outcomes (Pulcu & Browning, 2017). Research on depression suggests that the disorder may be associated with an asymmetry in these individual learning rates for rewarding and aversive stimuli. For example, patients with depression tend to change their choices rapidly after receiving one instance of negative feedback, even if on average their original choice results in positive outcomes (Murphy et al., 2003; Dombrovski et al., 2015). In addition to or in place of high learning rates for negative information, MDD patients may lack a bias towards learning from positive information, which is thought to be present in healthy adults. When healthy adults are asked to estimate the likelihood of an adverse event happening to them, such as a burglary, they update their initial estimates to a greater degree when they learn that the actual probability is lower than expected compared with when the true likelihood is higher than expected. Depressed individuals, on the other hand, show no such bias in updating their beliefs based on either positively or negatively valenced information (Garrett et al., 2014; Korn et al., 2014), although the severity of depressive symptoms appeared to be correlated with increased learning from negative feedback (Garrett et al., 2014).

Several studies modelling learning rates explicitly have supported the existence of a negative learning bias in depression. For instance, one study using a forced-choice decision making task indicated that people with depression had higher learning rates for negative feedback than healthy adults. In addition, learning rates were higher for negative than positive feedback within depressed participants (Cooper et al., 2014). Similarly, Beevers et al. (2013) found that depressed participants tended to have higher learning rates in punishment minimisation than reward maximisation conditions, suggesting increased responsiveness to negative outcomes. These findings are in accordance with classic cognitive models of depression, which propose that depression is characterised by a bias towards processing negative information compared to positive information (Beck, 2008).

However, opposite results have also been reported, with Dombrovski and colleagues (2010) indicating that suicide attempters had reduced learning rates for punishments but not rewards compared to healthy adults. A possible explanation of this unpredicted finding may be that depressed participants had a tendency to view negative outcomes as unchangeable or uncontrollable, and therefore did not alter their behaviour in reaction to such events. This might then elucidate potential causes for symptoms of learned helplessness of depression, which have been associated with perceived uncontrollability of adverse events (Seligman, 1972). There have been no additional studies confirming such a relationship, however, with some reporting either lower learning rates for both positive and negative outcomes (Chase et al., 2010), or no changes in learning rates associated with depression at all (Cavanagh et al., 2011; Gradin et al., 2011; Huys et al., 2013; Kunisato et al., 2012).

Considering the large variability in learning rates for depressed individuals during probabilistic learning tasks described here, it has been proposed that depression is more likely to be associated with a complex change in learning rates or strategies than a fundamental increase or decrease for negative and positive stimuli, respectively (Robinson & Chase, 2017). It has previously been emphasised that, in ecological environments, faster learning does not necessarily equate to better learning. Rather, the speed of learning should match the demands presented by the state of the environment (Scholl & Klein-Flügge, in press). One possibility is therefore that we need to search for failures in adapting learning rates to different environmental states to observe consistent learning changes associated with depression. For example, it has been suggested that depressed individuals may be able to initially acquire action-outcome associations, but have difficulties in adapting to alterations of the state of the environment. In a contingency-shifting version of the Iowa Gambling Task, participants with MDD tended to be less flexible in altering their choice

once they had developed a preference for a particular stimulus (Cella et al., 2010). Interestingly, however, depression severity was associated with increased learning from negative feedback, again pointing to enhanced learning rates for adverse information. Moreover, in reversal learning tasks, participants with MDD were found to be more likely to believe that a negative outcome indicated a reversal of action-outcome associations rather than a chance outcome compared to healthy controls (Taylor Tavares et al., 2008), further indicating increased weighting of negative information. These immediate reversals of choices following negative feedback were associated with a failure to recruit the ventrolateral prefrontal cortex and medial prefrontal cortex, providing some insight into potential underlying mechanisms for affected learning rates. Overall, these results indicate that individuals with depression might be unable to adjust learning rates accurately to positive outcomes in changing environments over time, whereas they are overly responsive to negative feedback in similar situations. These findings are not undisputed, however, as another study using a Probability Reversal Task did not observe either acquisition or reversal learning deficits in depressed individuals (Murphy, 2003). Additional research is therefore needed to clarify potential impairments in learning rate adaptations associated with depression.

One step forward in this particular topic might be to study meta-parameters associated with learning rates, such as the changeability of the environment. Previous studies have shown that, in healthy adults, learning rates tend to be higher in fast-changing or volatile environments than in stable environments (Behrens et al., 2007; Pulcu & Browning, 2017). Of interest, individuals with high levels of anxiety have shown smaller differences in learning rates for negative outcomes in stable and volatile conditions than adults with low anxiety levels (Browning et al., 2015), indicating that they were unable to adjust their learning behaviour in synchrony with alterations of the environment. Similar

impairments might be present in depressive disorders, which are closely related to symptoms of anxiety (Norman et al., 1996). In conclusion, the evidence to date suggests that depressed individuals may have a tendency to learn faster from negative than from positive information. Although differences in the learning rate might depend on the task demands, research studies to date have not explicitly addressed this hypothesis. While the exact changes in learning rates thus need to be explored further, they might present a potential behavioural target for interventions in depression. By reducing the weighting assigned to negative information when making choices, a balance in learning from different types of outcomes might be regained, which could consequently lead to reinforcement learning and decision making based on the entire spectrum of available information and/or positive sources of information in particular.

1.3.4. Learning strategy and model selection

Finally, changes in decision making performance may occur if depressed individuals adopt different learning strategies to guide their actions. Broadly speaking, reinforcement learning strategies can be classified as model-based or model-free, which apply different computations in order to optimise behaviour (Dayan & Niv, 2008). In model-free learning, individuals directly update their beliefs about action-outcome contingencies based on prediction errors (Chen et al., 2015). As such, this strategy is computationally undemanding, but is also inflexible and relatively unresponsive to sudden changes in the environment or in valuation of rewards (Adams, Huys, & Roiser, 2015). Model-based learning, in contrast, integrates experiences over time to construct an internal model of how actions can result in specific outcomes in a particular state. Using this strategy, a decision-maker can readily adapt to association reversals and other rapid changes within the environment (Dayan & Niv, 2008).

It has previously been suggested that model-free learning is unimpaired in

depression, such that automatic or reflexive responses to stimuli are intact (Beevers et al., 2013). Model-based learning, on the other hand, is thought to be impaired in people with depressive disorders (Huys, Daw, & Dayan, 2015; Rutledge et al., 2017). It can therefore be hypothesised that people with depression tend to default to using model-free rather than model-based learning strategies. In the Leapfrog Task discussed in section 1.3.2., for instance, data from depressed individuals fitted model-free learning computations better compared with results from non-depressed individuals (Blanco et al., 2013). This indicates that patients with depression may find it difficult to develop more extensive representations of the world, which can lead to suboptimal choice behaviour when navigating complex environments. This hypothesis of the presence of model-based learning deficits and its effects on reinforcement learning in depressed populations has not been tested thoroughly, however, and requires further exploration.

1.3.5. Overview of learning impairments in depression

Overall, the literature suggests that depression may be associated with alterations in sensitivity, choice stochasticity, learning rates, and model selection in reinforcement learning. While the evidence is highly heterogeneous across these parameters, the dominant view appears to be that depressed individuals are hyposensitive to rewards and hypersensitive to punishments, engage in more exploratory choice behaviour in reward-based than punishment-based conditions, have difficulties in adjusting their learning rates to changes in the environment, and fail to develop a complex understanding of action-outcome associations. It is at present unclear whether these deficits tend to occur concurrently or in isolation, and if these factors interact with one another. In addition, it is likely that there is heterogeneity in the specific learning patterns and their underlying neural mechanisms across individuals. Further research using computational models may be useful in identifying specific subtypes of Major Depressive Disorder associated with

different impairments in reinforcement learning. Overall, the findings suggest that each of the parameters discussed here may present a target of interest for interventions aiming to improve reinforcement learning impairments in depressive disorders.

1.4. Effects of NIBS on reinforcement-guided learning

Non-invasive brain stimulation can be applied to alter neural activity in specific brain regions, followed by testing for changes in information processing operations that are localized within those stimulated regions. It has previously been shown that reward and punishment processing engage a widespread cortico-striatal network (see Figure 1.4.), as well as frontal brain regions such as the anterior cingulate cortex (ACC) and the OFC (Admon & Pizzagalli, 2015). There is therefore a myriad of regions which could be targeted using NIBS techniques when attempting to alter reinforcement learning. Studies of learning and decision making, however, have tended to focus almost exclusively on the prefrontal cortex, in particular the DLPFC and the frontopolar cortex (FPC), which are associated with a wide range of higher-level processes including executive functioning and reasoning (Christoff & Gabrieli, 2000; Funahashi & Andreau, 2013). In the following section, I will provide an overview of NIBS studies carried out to date which have focused on modifying reinforcement learning processes in healthy and depressed populations (see Table 1.1.). In addition, I will briefly discuss potential novel targets for stimulation based on research investigating neural correlates of the individual reinforcement learning parameters.

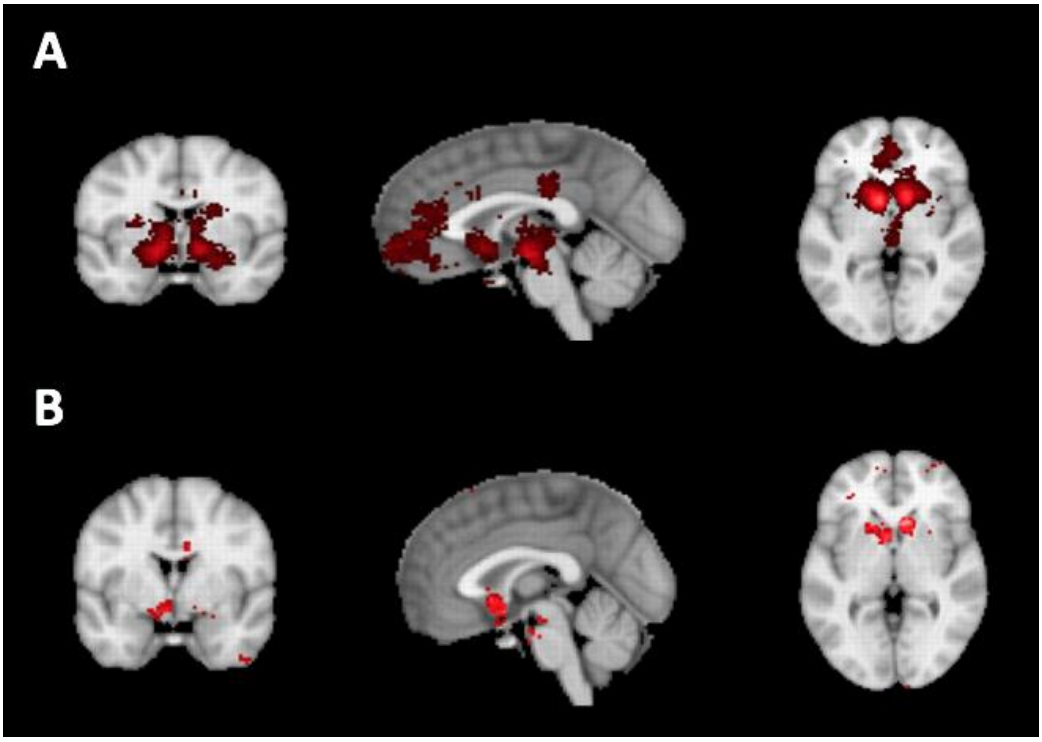

Figure 1.4. Meta-analytic results obtained from the Neurosynth database (http://neurosynth.org; Yarkoni et al., 2011). Each map represents foci reported in papers associated with specified key term. **(A)** Foci reported for the term "reward", based on 671 studies. Brain regions most commonly activated include the ventral and dorsal striatum, prefrontal cortex, insular cortex, thalami, and (para)cingulate gyrus. **(B)** Foci reported for the term "punishment", based on 63 studies. Brain regions most commonly activated include the ventral striatum, insular cortex, orbitofrontal cortex, and (para)cingulate gyrus.

1.4.1. Stimulation of the DLPFC alters reward sensitivity

In the context of pathological gambling and addiction, there has been much interest in altering risk-taking behaviour with the use of NIBS. As described in section 1.3.1., however, behaviour on risky decision making tasks may also provide insight into changes in reward sensitivity, and as such could therefore be of interest to research on depression. While there are no studies investigating alterations in risk-taking behaviour in depressed individuals with NIBS, there is evidence from healthy populations indicating that reward and punishment sensitivity may be changed by applying either TMS or tDCS.

A frequently employed task in these studies is the Balloon Analogue Risk Task

(BART; Lejuez et al., 2002), in which participants are asked to inflate a computerised balloon by pushing a virtual 'pump'. After each pump, participants have to decide to either continue inflating the balloon or to stop immediately. Participants accumulate money with each pump, while all gains are lost if the balloon is inflated too much and explodes. Importantly, each balloon has a different explosion point, so that participants are unable to estimate the optimal number of pumps across task trials. Fecteau and colleagues (2007b) showed that risk-taking behaviour on the BART tends to increase over time in healthy adults receiving sham tDCS, while tDCS targeting the right DLPFC (anode) and left DLPFC (cathode) reduced this change towards risky choices, indicating that participants may have become more sensitive to negative outcomes relative to rewards. Similar results were obtained if the anode was applied over the left DLPFC and the cathode was positioned over the right orbitofrontal cortex (Nejati et al., 2018). An MRI study, on the other hand, indicated that anodal tDCS of the right DLPFC did not result in altered decision making, although neural responses to losses were increased in the right DLPFC and ACC following stimulation (Weber et al., 2014). The authors suggest that the absence of a behavioural effect might be the result of using a reduced number of trials due to the necessity of using long durations of individual trials in the MRI scanner environment. Other studies have suggested that the effects of tDCS might depend on the particular population tested. In a task similar to the BART using card rather than balloon stimuli, anodal stimulation of the right DLPFC resulted in decreased risk-taking compared with sham stimulation in smokers, but increased risk-taking in non-smokers (Pripfl et al., 2013). In addition, several studies have reported no alterations in risk-taking behaviour on the BART with bifrontal tDCS (Cheng & Lee, 2016), even in a large sample size of 117 healthy adults (Russo et al., 2017). Responses to tDCS over the DLPFC as measured with BART performance thus appear to be variable depending on participant characteristics.

In addition to considering general risk-taking tendencies, a rough indication of reward sensitivity can be provided by measuring responses to varying reward magnitudes across trials. In the Risk Task, for instance, participants are presented with a number of boxes in two different colours, and are asked to pick a colour. Participants are awarded points for choosing the correct box colour containing a winning token, and are punished for choosing the incorrect colour. The ratio of the number of boxes in each colour is varied across the trials, and the amount of points which can be won by choosing each colour is explicitly presented. Participants are therefore required to weight the probability of a win for each colour (that is, the number of boxes) against the magnitude of the reward given for each colour. When bifrontal tDCS was applied at 2 mA targeting the right DLPFC, it was found that the influence of reward magnitude on choice behaviour was reduced in healthy participants (Fecteau et al., 2007a). That is, participants became more likely to choose safe options in the face of large potential rewards. This effect was not replicated in a sample consisting of older adults, however (Boggio et al., 2010). Combined with the findings on the BART, this suggests that anodal tDCS of the right DLPFC might increase sensitivity to punishments and reduce sensitivity to rewards, although the effects may be limited to healthy young adults.

Another task used to measure responsiveness to positive versus negative outcomes is the Weather Prediction Task, in which participants have to decide between one option which has a specified probability resulting in the presentation of an aversive image, and another option which can probabilistically lead to the presentation of an appealing image. Crucially, both choices are also associated with a reward which varies in magnitude across the trials. In certain trials, the risk of encountering a negative image might therefore be outweighed by the possibility of receiving a large reward. It has been shown that anodal tDCS over the right DLPFC reduced the number of choices which might result in a

negative outcome in individuals with high anxiety, but not in those with low anxiety (Chrysikou, Gorey, & Aupperle, 2017). Again, these findings suggest that the outcomes of tDCS on reward sensitivity might depend on particular characteristics or baseline trait markers of the individual.

One difficulty with the paradigms described above is that they do not allow for measuring changes in sensitivity to rewards or losses separately. Such a measurement would be of considerable interest, given that it has been proposed that there is an imbalance in sensitivity to rewards and punishments in depression. Several tasks have been developed which may provide further insight into whether this asymmetry might be malleable with the use of brain stimulation. Ye and colleagues (2015), for example, measured healthy adults' risk-taking behaviour in both gain- and loss-frames using a risk-measurement table. In this task, participants were presented with a table consisting of 16 two-alternative choices, each of which could result in different amounts of winnings in the gain-frame version, or different amounts of losses in the loss-frame version. Participants were explicitly shown the probability of the occurrence of each outcome, thereby enabling them to make an informed decision. When bifrontal tDCS was given with the anode over the right DLPFC at 2 mA, healthy participants were found to make riskier choices in the gain-frame, while they chose less risky options in the loss-frame. This effect of anodal stimulation over the right DLPFC was replicated in a follow-up study using similar risk-measurement tables (Ye et al., 2016). In addition, when the anodal electrode was positioned over the left DLPFC, participants similarly showed a greater number of risky choices in the gain frame, although contrary to right-sided stimulation no effects were observed in the loss-frame (Ye et al., 2016). In contrast with the majority of studies investigating responses on the BART, the Risk Task, and the Weather Prediction Task, these findings suggest that tDCS targeting the right DLPFC might be able to increase

sensitivity to rewards and reduce sensitivity to adverse outcomes in healthy adults.

Inferring reward sensitivity from risky behaviour, however, is not without its problems. Firstly, behaviour might be the result of other factors such as impulsivity as well as changes in sensitivity. In addition, these tasks do not provide a great deal of insight into how individuals adapt their behaviour over time as a result of their experiences. Only one study to date has specifically investigated loss and reward sensitivity in a probabilistic learning context, where participants were required to learn which one of two paired symbols resulted in a better outcome (Ott et al., 2011). In this study, continuous theta-burst stimulation (cTBS), a TMS protocol that has been shown to inhibit cortical excitability when applied to motor cortex, was applied over the left DLPFC. It was found that this procedure enhanced sensitivity to wins during learning of action-outcome associations as indicated by increases in approach learning. If the stimulation was targeting the right DLPFC, on the other hand, participants tended to become more sensitive to negative outcomes when making their decisions.

Across these studies, there is thus some evidence that increasing activity in the right DLPFC can enhance sensitivity to negative outcomes and potentially reduce sensitivity to positive outcomes. Furthermore, there are indications that increased activity in the left DLPFC may change sensitivity to positive outcomes, although fewer studies have investigated this possibility. The effects of stimulation are complex, as they appear to be dependent not only on the tasks used to investigate sensitivity but also on differences in population characteristics. Further complications arise when taking into account that activity in reward-sensitive areas throughout the brain are highly context-dependent. For instance, it has been suggested that a gain might be processed as a negative outcome if it is lower than other potential outcomes (Nieuwenhuis et al., 2005). In some of the tasks previously described, which require participants to choose between two positive or two

negative options, teasing apart behavioural changes associated with either reward- or punishment-based sensitivity can therefore become convoluted. Furthermore, as previously stressed, the majority of the tasks described here are focused on examining risky behaviour rather than sensitivity to positive or negative stimuli, and are suboptimal for examining how brain stimulation might affect this parameter. Novel paradigms may therefore need to be developed in order to gain a better understanding of the effects of tDCS and TMS on sensitivity. In addition, improved differentiation of sensitivity changes from other parameters driving choice behaviour might be achieved by explicitly including this parameter in mathematical models (see Robinson & Chase, 2017 for an example).

Given the hints that results might differ depending on trait and state characteristics of the participants, further research is also needed to determine whether the effects of stimulation in depressed individuals are comparable to those observed in healthy adults. In addition, alternative brain regions might be of interest for studies looking at brain stimulation in the field of reward and punishment sensitivity instead of just the DLPFC. Activation in the left and right ventral striatum, for example, has previously been associated with encoding the value of rewards (McGuire et al., 2014), and was found to correlate with healthy individuals' sensitivity to rewards (Kim et al., 2015). In addition, BOLD signals in the ventromedial prefrontal cortex (VMPFC) and the orbitofrontal cortex have been correlated with reward values of choices (O'Doherty, 2004; Rushworth et al., 2011). Sensitivity to punishment, on the other hand, has been associated with activation in the right lateral OFC during feedback and the left dorsal striatum in anticipatory stages of aversive outcomes (Kim et al., 2015). These differences in neural mechanisms associated with processing reward and punishment outcomes present the possibility that sensitivity to different outcomes might be altered independently.

This hypothesis could be explored further by targeting associated regions with non-

invasive brain stimulation. While direct effects of NIBS are typically limited to the cortical surface, it has been demonstrated that it is also possible to stimulate deeper structures such as the striatum indirectly with both TMS and tDCS by targeting connected brain regions. For example, dopamine release and reward-related processing in the dorsal striatum have been found to be altered with the application of rTMS to the primary motor cortex (Strafella et al., 2003) and the anterior prefrontal cortex (Van Holstein et al., 2018), respectively. Similarly, it has been shown that anodal tDCS of the DLPFC can result in widespread changes in perfusion in anatomically connected brain regions such as the thalami (Stagg et al., 2013). In addition, adaptations of NIBS techniques, such as novel coil designs in TMS (Crowther et al., 2011; Zangen et al., 2005) have been developed to improve stimulation of subcortical brain regions. Future research may therefore utilise TMS and tDCS to assess whether modulation of reward sensitivity is possible through changing ventromedial, orbitofrontal, or striatal activation.

1.4.2. NIBS can affect deliberate exploration

In addition to changing reward sensitivity, stimulation of the prefrontal cortex might result in alterations of exploratory or stochastic decision making in reinforcement learning. Several studies have targeted the frontopolar cortex when investigating this possibility, as this region has been implicated in driving exploratory behaviour (Cavanagh et al., 2012; Domenech & Koechlin, 2015; Laureiro-Martínez et al., 2014).

Of significance, a study investigating cTBS over the right frontopolar cortex has indicated that the neural networks underlying stochastic and deliberative decision making are distinct (Zajkowski, Kossut, & Wilson, 2017). Using the 'Horizon Task', a bandit task which was specifically designed to separate randomness and deliberate exploration, it was found that cTBS reduced deliberate but not random exploration. With tDCS, however, stimulation of the same regions resulted in increased exploratory behaviour on a virtual

slot machine task (Raja Beharelle et al., 2015). Here, participants were required to choose among three bandits which could result in different reward magnitudes. It was found that anodal stimulation of the right frontopolar cortex with 1 mA tDCS reduced inverse temperature values. The authors propose that this change signals an increase in deliberate exploration rather than choice stochasticity, as the altered values were associated with increased reaction time, thereby indicating greater consideration of the available options. Although both the tDCS and cTBS protocols used in these two studies aimed to increase neural excitability, it is possible that the disparate findings might have resulted from differences in the underlying physiological mechanisms for 'offline' (stimulated at rest) TMS and 'online' (stimulated during task performance) tDCS. In addition, tDCS tends to result in less focal stimulation than TMS (Sparing & Mottaghy, 2008), and may therefore inadvertently reach additional frontal regions. Alternatively, the task used by Zajkowski and colleagues (2017) was specifically designed to differentiate between random and deliberate exploration, and thus may have been more suitable for investigating these distinct processes compared to the task used by Raja Beharelle et al. (2015). It would therefore be of interest to directly compare the effects of tDCS and TMS within this computational task (Zajkowski et al., 2017) in order to gain a better understanding of the differential effects of the two NIBS techniques on choice stochasticity.

In addition to investigations of stimulation of the frontopolar cortex, one study has looked at effects of tDCS over the DLPFC on stochastic decision making. In a Probabilistic Learning Task (similarly employed by Ott et al., 2011), participants were presented with three pairs of symbols which were probabilistically associated with reward outcomes. In this paradigm, participants have to learn which of the two choices in each pair results in the best outcome. Following a training phase, participants are randomly presented with all possible combinations of the symbols in a transfer phase. In this task, it was found that

anodal tDCS of the left DLPFC with the cathode over the temporal cortex was associated with reduced inverse temperature values at the end of the training phase (Turi et al., 2015). In other words, stimulation of the DLPFC resulted in reduced correspondence between the estimated action values and choice behaviour during implicit probabilistic learning.

While the prefrontal cortex is a strong candidate for modulating exploration and exploitation behaviour, this parameter is thought to be affected by an extensive neural network and may thus be altered through stimulation of other brain regions. It has been suggested that phasic and tonic discharges of locus coeruleus (LC) neurons correlate with exploitation and exploration choices, respectively (Usher et al., 1999). Activity in these neurons is proposed to be regulated by the anterior cingulate cortex and orbitofrontal cortex (Aston-Jones & Cohen, 2005). In transgenic rats, it has been found that perturbing activation in the ACC through muscimol injections, an agonist of the inhibitory neurotransmitter GABA-A, led to increased stochasticity of decisions (Tervo et al., 2014). The ACC may therefore present an alternative target for neurostimulation studies investigating exploration and randomness of choices. To date, two studies have attempted to alter activation in this region by stimulating the DLPFC, which has strong anatomical connections to the dorsal anterior cingulate cortex, using both TMS (Paus, Castro-Alamancos, & Petrides, 2001) and tDCS (Weber et al., 2014).

Additional regions which have been implicated in exploratory behaviour include the ventral tegmental area and basolateral amygdala (Ishii, Yoshida, & Yoshimoto, 2002). A potential target for stimulation in this context might be the ventromedial prefrontal cortex (VMPFC), which receives input from these regions (Maroun, 2006; Walsh & Han, 2014). Previously, it has been shown that GABA and glutamate concentrations in the VMPFC predict average choice accuracy in a decision making task, with high levels of GABA and low levels of glutamate predicting high inverse temperature values in male

adults (Jocham et al., 2014). tDCS targeting the VMPFC was found to increase choice stochasticity in a decision making task in which the values and probabilities of different outcomes were explicitly presented (Hämmerer et al., 2016), which can be hypothesised to be caused by changes in GABA-ergic or glutamatergic changes in this region. It thus appears to be possible to affect exploratory and/or stochastic decision making using tDCS and TMS, although the large variety of stimulation parameters and tasks used make it difficult to identify potentially consistent effects across studies. Furthermore, the separation of stochastic and deliberate exploratory behaviour is currently not standard practice, which may dilute findings from NIBS studies tapping into these separate processes. Although further research is thus needed to clearly distinguish between effects on random and deliberate exploration and understanding which stimulation parameters are most effective in altering inverse temperature values, the findings reported here are promising in that this parameter appears to be susceptible to alteration with tDCS and TMS. Importantly, it is currently unclear whether such behaviours can be altered independently for positive and negative action-outcome associations. This will be of particular interest in depression-related research, as prior studies indicated that stochastic choices may be especially high in reward learning compared with punishment learning in individuals with depression (see section 1.3.2.).

1.4.3. Excitation of the left DLPFC increases learning rates in healthy adults

In probabilistic learning paradigms, there has been much interest in the effects of NIBS on learning rates. Prefrontal regions in particular have been investigated frequently, as it is known that the prefrontal cortex plays a key role in higher-order cognitive functions related to reward processing. In primates, for example, it has been shown that the DLPFC uses information about reward values to control behaviour (Wallis & Miller, 2003), while in humans activity in both the lateral and orbital prefrontal cortices has been related to the

ability to adjust actions in response to novel information (Ridderinkhof et al., 2004).

Using the Weather Prediction Task described in section 1.4.1, it has been found that offline tDCS at 2 mA with the anode over the left DLPFC and the cathode over the right DLPFC reduced implicit learning compared to sham tDCS in patients with MDD (Brunoni et al., 2013c). In contrast, in an earlier study, when the anode was applied over Fp3, which is located near the left DLPFC, and the cathode was placed over the vertex, learning was improved in healthy adults (Kincses et al., 2003). Differences between the study outcomes might be caused by the different populations and protocols investigated. Alternatively, differences in duration of the stimulation may have led to diverging results. For example, a prior study has reported that anodal tDCS at 2 mA over the primary motor cortex, which has excitatory after-effects when applied for 13 minutes, resulted in inhibitory after-effects when applied for 26 minutes (Monte-Silva et al., 2013). This suggests that the polarity of tDCS effects may reverse when applied for longer durations. While it is unclear whether such effects also occur with alternative electrode montages targeting cognitive rather than motor functions, it is possible that the 30-minute application of anodal tDCS in the study by Brunoni et al. (2013c) had similar reversed effects on excitation in the DLPFC. Hypothetically, inhibition of this region could accordingly result in decreased learning on the task. Finally, a potential explanation may be that cathodal stimulation over the right DLPFC reduced activation in this region, which has been found to influence learning and decision making processes (Xue et al., 2013; Nihonsugi, Ihara, & Haruno, 2015). Thus, it can be speculated that stimulation of only the left prefrontal cortex might lead to enhanced learning compared with bifrontal tDCS. The effects of stimulation of the left DLPFC on learning rates have been further explored in a tDCS study employing a value-based decision making task in which participants had to choose between one of two stimuli (Hämmerer et al., 2016). In contrast with the Weather

Prediction Task, the amount and probability of reward for each option were explicitly presented. In this task, anodal tDCS of the left DLPFC with the cathode positioned over the occipital region below the inion was associated with increased learning for rewarding stimuli compared to sham stimulation. Anodal stimulation of the left dorsolateral prefrontal cortex may thus enhance learning rates in both implicit and explicit reward learning tasks.

When utilising the Probabilistic Learning Task (see section 1.4.1.), however, there is little evidence that modulation of prefrontal cortical activity results in changes of learning rates with either TMS or tDCS. cTBS of the left DLPFC, for instance, did not alter overall learning efficacy on this task (Ott et al., 2011). Similarly, tDCS of the frontopolar cortex (Turi et al., 2015) or the anterior prefrontal cortex (Casula et al., 2017) did not affect learning performance in healthy adults. One possibility is that learning rates are more readily modulated in tasks which explicitly show the probability and values of different options, rather than relying on participants acquiring this knowledge implicitly. However, using the implicit Iowa Gambling Task, He and colleagues (2016) found that high-definition tDCS with anodal stimulation of the left DLPFC altered learning compared to sham stimulation. While learning rates were not modelled computationally, active tDCS was associated with faster learning than sham tDCS, implying an increase in learning rates. In addition, in an implicit reversal learning task, left DLPFC transcranial alternating current stimulation (tACS), which affects brain oscillations, induced faster learning following reversals of action-outcome associations (Wischnewski, Zerr, & Schutter, 2016). This behavioural change was associated with a decrease in theta-beta EEG ratio in the left hemisphere. In addition to prefrontal regions, there may be a role of the temporoparietal junction (TPJ) in reward learning. It has been suggested that the right TPJ affects the updating of expectations of the environment, which in turn leads to optimisation

of responses (Geng & Vossel, 2013). With double-pulse TMS, healthy individuals were found to have decreased learning rates in a probabilistic learning task if the stimulation was applied over the right TPJ 300 ms after the presentation of the target (Mengotti et al., 2017).

Overall, the evidence indicates that excitation of the left DLPFC can increase learning rates, with one study suggesting that inhibition of the right TPJ might reduce learning rates. In future NIBS studies, it may also be worth investigating the role of alternative brain regions in learning rates. A potential candidate is the orbitofrontal cortex (OFC), which has been suggested to be involved in encoding and updating utilities of different stimuli (Lee et al., 2007), with lesions of the OFC inhibiting flexible learning (Sadacca, Wikenheiser, & Schoenbaum 2017; Tsuchida, Doll, & Fellows, 2010). Activity of the dorsomedial frontal cortex (DMFC) has also been associated with learning rates in humans (McGuire et al., 2014). In rats, damage to the medial prefrontal cortex was associated with lower learning rates in reinforcement paradigms (Laskowski et al., 2016). Closely related to DMFC activity, it has been suggested that the ACC activity is particularly concerned with associations between rewards and actions (Rushworth et al., 2011) and learning rate adaptation (Behrens et al., 2007). The possibility of modifying activity within this region to change learning rates will be further explored in this thesis (see Section 1.7.).

1.4.4. Preliminary evidence for disruption of model-based learning with TMS of the DLPFC

At the highest level, NIBS might be able to modulate *in what way* participants learn from reinforcing stimuli, rather than determining *how much* they learn. That is, neurostimulation techniques might alter model-based and/or model-free learning mechanisms. In neuroimaging studies, model-free and model-based strategies have been found to be

associated with different neural pathways. Model-based error signals are mainly observed in a frontoparietal network involving the lateral prefrontal cortex, posterior intraparietal sulcus, and ventromedial prefrontal cortex. In addition, several studies have suggested a strong involvement of the medial orbitofrontal cortex in model-based learning in humans and animals (McDannald et al., 2011; Valentin, Dickinson, & O'Doherty, 2007). In contrast, model-free errors are generally correlated with activation in the ventral and dorsolateral striatum (Doll, Simon, & Daw, 2012; Gläscher et al., 2010; Maia, 2009; Wan Lee, Shimojo, & O'Doherty, 2013). One study, however, found that ventral striatum activity in rats was also associated with model-based learning (McDannald et al., 2011), indicating that the two networks may not be entirely independent. Rather than considering the two systems in isolation, some studies have attempted to break down the arbitration process for selecting model-based versus model-free learning strategies. Wan Lee and colleagues (2013), for instance, proposed that anterior lateral and polar prefrontal cortices are involved in computing estimates of the reliability of the two respective strategies, aiding in the choice of which policy is most appropriate. By targeting prefrontal cortex, NIBS techniques could be used either to attempt to alter the ability to apply model-free versus model-based learning, or to potentially influence the arbitration process of choosing one tactic over the other.

The few NIBS studies examining model-based learning as yet have focused exclusively on modulating DLPFC activity. Smittenaar and colleagues (2013) looked at the effects of cTBS on model-based learning using a two-step Markov task design. In this paradigm, each trial consists of two choice stages in which participants are presented with a two-alternative forced choice (Daw et al., 2011). At the first stage, the same two stimuli are presented on all trials, with each of the options being probabilistically associated with the presentation of one of two second-stage stimulus pairs. At the second stage, the four

different stimuli have a reward probability which varies between 0.2 and 0.8. Using a model-based strategy, participants should be able to learn which of the first-stage stimuli leads to the optimal second-stage pair more frequently. With model-free learning, however, received rewards would directly increase the likelihood of choosing the same stimuli on the next trial. When cTBS was applied over the right DLPFC, model-based learning was disrupted in healthy participants, whereas model-free learning was unaffected. This suggests that the right DLPFC plays a key function in the use of model-based learning mechanisms. Similarly, application of TMS at 1 Hz over the left LPFC was found to impair performance on a 3-stage Markov task (Wittkuhn et al., 2018), which is suggestive of an impairment in using model-based strategies. Disrupting excitability of the left lateral prefrontal cortex therefore appears to impair model-based learning. When the right DLPFC was targeted with tDCS, on the other hand, no changes in either model-based or model-free learning were observed compared to sham stimulation (Smittenaar et al., 2014). There is thus no evidence to date that tDCS can modulate the model strategy used by individuals during reward learning. However, definite conclusions about the relative efficacy of TMS and tDCS procedures cannot be drawn at this stage due to the sparsity of research investigating the effects of NIBS on either model-based or model-free learning.

Table 1.1. Stimulation parameters and results of studies assessing the impact of tDCS and TMS on reinforcement learning components

Parameter	First author (year)	*N*	Stimulation type	Design	Target region	Stimulation parameters	Control condition	Task	Main finding
Sensitivity	Boggio (2010)	28	Online tDCS	Between-subjects	Left DLPFC (anode) and right DLPFC (cathode)	2 mA for 15 minutes	Reverse electrode montage or sham tDCS	Risk Task	In older adults, stimulation did not affect task performance
	Cheng (2016)	16	Online tDCS	Within-subjects	Left DLPFC (anode) and right DLPFC (cathode)	2 mA for max. 19 minutes	Sham tDCS	BART	tDCS did not affect task performance
	Chrysikou (2017)	63	Online tDCS	Between-subjects	Right DLPFC (anode) and contralateral mastoid (cathode)	1.5 mA for max. 20 minutes	Stimulation of right DLPFC or sham tDCS	Weather Prediction Task	tDCS decreased risky choices in individuals with high anxiety, but did not change behaviour of low-anxious participants
	Fecteau (2007a)	35	Online tDCS	Between-subjects	Left DLPFC (anode) and right DLPFC (cathode)	2 mA for max. 20 min	Reverse electrode montage or sham tDCS	Risk Task	With tDCS, the influence of reward magnitude on choice behaviour was reduced
	Fecteau (2007b)	36	Online tDCS	Between-subjects	Right DLPFC (anode) and left DLPFC (cathode)	2 mA for max. 15 min	Reverse electrode montage or sham tDCS	BART	Active tDCS increased sensitivity to negative outcomes relative to positive outcomes
	Nejati (2018)	24	Offline tDCS	Within-subjects	Left DLPFC (anode) and right OFC (cathode)	1.5 mA for 20 min	Reverse electrode montage or sham tDCS	BART	Active tDCS decreased risk-taking, indicating greater sensitivity to negative outcomes relative to positive outcomes

Table 1.1. *(Continued)*

Parameter	First author (year)	*N*	Stimulation type	Design	Target region	Stimulation parameters	Control condition	Task	Main finding
Sensitivity	Ott (2011)	42	cTBS	Between-subjects	Left DLPFC or right DLPFC	600 pulses at 50 Hz	Stimulation of right DLPFC or sham tDCS	Probabilistic Learning Task	TMS of the left DLPFC increased reward sensitivity, while stimulation of the right DLPFC increased punishment sensitivity
	Pripfl (2013)	36	Online tDCS	Within-subjects	Left DLPFC (anodes) and right DLPFC (cathodes)	0.45 mA for 15 min	Reverse electrode montage or sham tDCS	Columbia Card Game	tDCS increased risk-taking in non-smokers, but had the opposite effect in smokers
	Russo	198	Online tDCS	Between-subjects	Right DLPFC (anode) and left DLPFC or supraorbital area (cathode)	2 mA for 30 min	Reverse electrode montage or sham tDCS	BART	tDCS did not affect task performance
	Weber (2014)	22	Offline tDCS	Between-subjects	Right DLPFC (anode) and left DLPFC (cathode)	1.5 mA for 15 min	Sham tDCS	BART	Active tDCS did not alter behaviour, but increased activity in the right DLPFC and ACC
	Ye (2015)	60	Offline tDCS	Between-subjects	Right DLPFC (anode) and left DLPFC (cathode)	2 mA for 15 min	Reverse electrode montage or sham tDCS	Risk Measurement Table	tDCS increased risky behaviour in gain-frames, whereas it decreased risk-taking in loss-frames

Table 1.1. *(Continued)*

Parameter	First author (year)	*N*	Stimulation type	Design	Target region	Stimulation parameters	Control condition	Task	Main finding
Sensitivity	Ye (2016)	100	Offline tDCS	Between-subjects	Right DLPFC or left DLPFC (anode) and Pz (cathode)	2 mA for 15 min	Anodal tDCS of the left DLPFC, cathodal tDCS of the right or left DLPFC, or sham tDCS	Risk Measurement Table	tDCS of the right or left DLPFC increased risk-taking in gain frames. Only right-sided tDCS reduced risky choices in the loss-frame
Stochasticity	Hämmerer (2016)	16	Online tDCS	Within-subjects	VMPFC (anode) and occipital cortex (cathode)	2 mA for 15 min	Anodal tDCS of lateral PFC and sham tDCS	Value-based Decision Making Task	Stimulation of the VMPFC was correlated with increased choice stochasticity
	Raja Beharelle (2015)	79	Online tDCS	Between-subjects	Right FPC (anode) and vertex (cathode)	1 mA, timing unclear	Cathodal or sham tDCS	Virtual Slot Machine Game	Stimulation was associated with increases in deliberate exploratory behaviour
	Turi (2015)	16	Online tDCS	Within-subjects	Left DLPFC (anode) and temporal cortex (cathode)	1 mA for 16 min	Sham tDCS	Probabilistic Learning Task	tDCS increased stochastic behaviour during reinforcement learning
	Zajkowski (2017)	26	cTBS	Within-subjects	Right FPC	600 pulses at 50 Hz	Stimulation of the vertex	Horizon Task	Stimulation reduced deliberate but not random exploration in healthy adults

Table 1.1. *(Continued)*

Parameter	First author (year)	*N*	Stimulation type	Design	Target region	Stimulation parameters	Control condition	Task	Main finding
Learning rate	Brunoni (2013c)	28	Online tDCS	Between-subjects	Left DLPFC (anode) and right DLPFC (cathode)	2 mA for 30 min	Sham tDCS	Weather Prediction Task	In patients with MDD, tDCS impaired implicit learning performance
	Casula (2017)	18	Online tDCS	Within-subjects	Anterior PFC (anode) and occipital cortex (cathode)	2 mA for 25 min	Sham tDCS	Probabilistic Learning Task	tDCS did not affect learning performance in healthy adults
	Hämmerer (2016)	16	Online tDCS	Within-subjects	Left DLPFC (anode) and occipital cortex (cathode)	2 mA for 15 min	Anodal tDCS of lateral PFC and sham tDCS	Value-based Decision Making Task	Learning rates were increased in healthy adults with anodal stimulation
	He (2016)	110	Offline high-definition tDCS	Between-subjects	Left DLPFC (anode) and F5, AF3, FC3, and F1 (cathodes)	1.5 mA for 20 min	Sham tDCS	Iowa Gambling Task	tDCS was related to increased learning rates in healthy adults
	Kincses (2003)	22	Online tDCS	Within-subjects	PFC (anode) and vertex (cathode)	1 mA for 10 min	tDCS of the occipital cortex or sham tDCS	Weather Prediction Task	Stimulation enhanced learning in healthy adults
	Mengotti (2017)	21	Online TMS	Within-subjects	Right TPJ	Double-pulse TMS	Sham tDCS	Location-cueing paradigm	TMS impaired learning rates when applied 300 ms after stimulus presentation

Table 1.1. *(Continued)*

Parameter	First author (year)	*N*	Stimulation type	Design	Target region	Stimulation parameters	Control condition	Task	Main finding
Learning rate	Ott (2011)	42	cTBS	Between-subjects	Left DLPFC	600 pulses at 50 Hz	Stimulation of right DLPFC or sham tDCS	Probabilistic Learning Task	There were no changes in learning performance observed with stimulation
	Turi (2015)	16	Online tDCS	Within-subjects	Left DLPFC (anode) and temporal cortex (cathode)	1 mA for 16 min	Sham tDCS	Probabilistic Learning Task	tDCS did not affect learning in healthy adults
	Wishnewski (2016)	50	Online tACS	Between-subjects	Left (anode) and right (cathode) frontal cortex	1 mA for approximately 11 min	Sham tACS	Reversal Learning Task	tACS induced more rapid learning of action-outcome association reversals
Model selection	Smittenaar (2013)	25	Offline cTBS	Within-subjects	Right DLPFC	600 pulses at 5 Hz	cTBS to the left DLPFC and vertex	Markov Decision Task	TMS impaired participants' ability to engage in model-based learning
	Smittenaar (2014)	22	Online tDCS	Within-subjects	Right DLPFC (anode) and inion (cathode)	2 mA for 25 min	Sham tDCS	Markov Decision Task	There were no changes in model-based learning with tDCS
	Wittkuhn (2018)	31	TMS	Within-subjects	Left LPFC	1200 pulses at 1 Hz	TMS of the vertex	Markov Decision Task	TMS reduced performance on the model-based learning task

Abbreviations: BART = Balloon Analogue Risk Task; cTBS = Continuous Theta Burst Stimulation; DLPFC = dorsolateral prefrontal cortex; FPC = frontopolar cortex; LPFC = lateral prefrontal cortex; PFC = prefrontal cortex; OFC = orbitofrontal cortex; tACS = Transcranial Alternating Current Stimulation; tDCS = Transcranial Direct Current Stimulation; TMS = Transcranial Magnetic Stimulation; TPJ = temporoparietal junction; VMPFC = ventromedial prefrontal cortex.

1.5. Discussion

1.5.1. Overview of the findings

The ability to apply reinforcement learning is crucial to making optimal decisions in daily life, and may be impaired at various stages in depressive disorders. In this review, I have proposed that these deficits may be separated into discrete processes using computational modelling. Moreover, there is potential for the use of non-invasive brain stimulation techniques such as tDCS and TMS for ameliorating such impairments.

Firstly, reinforcement learning might be altered due to changes in sensitivity to rewards and punishments. Overall, the evidence suggests that individuals with depression are hyposensitive to rewards and hypersensitive to negative outcomes. In healthy adults, sensitivity might be modified by stimulating the left or right DLPFC, although the direction of effects on sensitivity tends to vary across tasks and studies. Secondly, depression has been related to increased random or exploratory behaviour in gain-contexts, while such behaviour was found to be reduced in loss-contexts. However, it should be noted that these valence-specific changes in exploratory behaviour could be due to underlying differences in reward sensitivity. Additional research is needed to establish whether choice stochasticity and reward sensitivity have separate effects on reward learning in depression. With tDCS and TMS over prefrontal regions, choice stochasticity and/or deliberate exploration have been altered in healthy participants. However, due to the great variety in tasks and stimulation parameters used in the studies conducted to date, it is presently unclear if NIBS can reliably alter random or deliberate exploratory behaviour in a specific direction. Moreover, it is unknown whether these factors can be altered specifically for either reward or punishment outcomes. Thirdly, there are indications that depressed individuals might have altered learning rates for positive and negative stimuli compared to healthy controls, although findings have been heterogeneous

and may depend on task demands. In NIBS studies, stimulation of the left DLPFC was found to increase learning rates, while there is preliminary evidence from one study suggesting that inhibition of the right TPJ might reduce learning. Finally, it has been proposed that model-based learning is impaired in depression, whilst model-free learning is unaffected. Stimulation of either the right or the left DLPFC has been found to reduce model-based learning with TMS in two studies, whereas no significant effects have been observed with the use of tDCS. Further studies are needed to explore whether model-based learning might also be enhanced rather than perturbed using non-invasive brain stimulation.

1.5.2. Future directions

It thus appears that modulation of neural activity, particularly in prefrontal regions, can result in changes of reinforcement learning processes associated with depression. One complication in interpreting the findings from research carried out to date is that the majority of task paradigms and/or analyses do not allow for clear separation of the individual parameters outlined in computational models. Increased use of computational modelling in NIBS studies may enhance our ability to differentiate between changes in different processes involved in solving cognitive tasks. While there have been several studies applying such models in neurostimulation designs, this is a relatively novel development, and there are therefore certain caveats in the current evidence base which deserve further exploration. Firstly, research has almost exclusively focused on healthy populations, with only one of the 28 studies discussed here examining effects of tDCS in depressed individuals. As alterations in reinforcement learning might be of particular value in the development of treatments for psychiatric disorders, it is essential to investigate whether the findings discussed here translate to people with depression and related disorders. In addition, little is known about the effects of either TMS or tDCS at a neural

level in decision making tasks. As the networks involved in reward and punishment learning are extensive and NIBS can have effects not only in the targeted region but also on distal brain areas, it is unknown which underlying changes in neural activity are responsible for the observed alterations in behavioural performance. Future studies might contribute to our understanding of NIBS effects by applying neuroimaging methods to assess changes in decision making as a consequence of stimulation. As computational modelling can also be applied to neuroimaging data (Gläscher & O'Doherty, 2010), it is possible to learn how neural representations of specific computational parameters are altered with tDCS or TMS.

1.5.3. Research aims and thesis outline

In conclusion, NIBS may be useful for modulating reinforcement learning, and can potentially affect individual components of the learning process. Further research informed by computational theory may assist in clarifying the mixed results reported to date and isolating the contribution of closely related constructs such as reward sensitivity, exploratory behaviour, and speed of learning. In this thesis, I propose that impairments in learning rates are of particular interest in the development and treatment of depression. Previously, it has been hypothesised that the learning rate reflects people's estimate of the *information content* of events, or how much an event can improve predictions of future outcomes (Behrens et al., 2007; Browning et al., 2015; Pulcu & Browning, 2017). In other words, if an event is thought to provide a great deal of information about an action-outcome association (i.e., has a high information content), people should update their expectations to a large extent and will therefore have high learning rates. Events that provide little information about the action-outcome association (i.e., have a low information content), on the other hand, should influence people's predictions to a lesser degree and result in low learning rates.

In order to optimise their behaviour, individuals are expected to be biased towards processing events that they believe have the highest information content (Behrens et al., 2007; Browning et al., 2015; MacKay, 2003; Pulcu & Browning, 2017). Therefore, if an individual believes that negative events are more informative than positive events when trying to predict future outcomes, they may develop an affective bias towards attending to adverse events. Underlying beliefs of how informative negative events are could thus explain the presence of negative attention biases characteristic of depressive disorders (Pulcu & Browning, 2017). Based on this theory, it can be predicted that adjusting people's beliefs about the information content of positive and negative events as reflected by learning rates can change associated affective biases, which in turn may ameliorate symptoms of depression.

Importantly, the present review proposed that depression may be related with more complex alterations of learning rates than previously thought. For example, rather than presenting with global increases in negative learning rates and/or reductions in positive learning rates, depressed individuals may have specific difficulties in adjusting their speed of learning to changes in statistical aspects of the environment such as volatility (see section 1.3.3.). In this thesis, I will therefore focus on examining whether learning rates can be modified in volatile and stable environments with the application of transcranial direct current stimulation in a series of experimental studies. In Chapter 2, I discuss the theoretical and methodological principles guiding the design of the experimental research studies. In Chapter 3, I provide an overview of the test-retest reliability and practice effects on the Information Bias Learning Task, the main paradigm used to assess effects of neurostimulation on reinforcement learning in this thesis. In Chapter 4, I examine the impact of tDCS targeting the dorsolateral prefrontal cortex on learning rates and other computational components in healthy adults. In Chapter 5, the influence of concurrent

tDCS and a training version of the IBLT on learning parameters and affective biases in healthy adults is investigated. Finally, in Chapter 6, I analyse the effects of prefrontal tDCS on learning in a pilot study of individuals with subclinical depression. The main aim of this work is to contribute to our understanding of how non-invasive brain stimulation may be applied to alter reinforcement learning processes, and thereby advance research on potential novel interventions for depression.

Chapter 2. Methods for stimulating neural networks of reinforcement learning

2.1. Introduction

In this thesis, I will examine whether learning rates and consequent affective biases can be modified with the application of transcranial direct current stimulation (tDCS). The present chapter provides a description of the main cognitive task and non-invasive brain stimulation technique used to alter reward learning rates in Chapters 4-6. First, I will discuss the overarching logic of the Information Bias Learning Task (IBLT), a computerised reinforcement learning task in which the information content and consequent learning rates for positive and negative outcomes can be manipulated. Next, I will summarise the body of work investigating the physical principles and neurophysiological mechanisms of tDCS. Finally, I will describe the rationale for the tDCS protocol used to alter performance on the Information Bias Learning Task in the neurostimulation studies reported here.

2.2. Modulating the information content of events

2.2.1. Expected and unexpected uncertainty

Computational work has demonstrated that one way of altering estimations of the information content of events, as reflected in learning rates, is to modify sources of uncertainty about the relationship between actions and outcomes (Behrens et al., 2007; Browning et al., 2015; Pulcu & Browning, 2017). A first cause of uncertainty is what is generally known as expected uncertainty, in which a choice results in a specified outcome on some but not all occasions (Payzan-LeNestour et al., 2013). For example, if a choice is probabilistically associated with a positive outcome 70% of the time, a relatively surprising outcome will follow the choice approximately 30% of the time due to chance. When outcomes have a low expected uncertainty, and are therefore highly informative

about the underlying association (as in this 70-30% example), learning rates should be higher than in an environment where the expected uncertainty is high (e.g. a 50-50% probability).

A second influential uncertainty factor is the volatility of the environment, a form of *unexpected* uncertainty which indicates how changeable probabilistic relationships are over time (Yu & Dayan, 2005). In highly stable conditions, an unexpected outcome is likely to be due to noise or expected uncertainty rather than to an alteration in the underlying probabilities. In this scenario, to successfully predict future outcomes, individuals should take into account their experiences over an extended period of time rather than basing their action choices mainly on recent surprising outcomes. That is, if the environmental state is stable, each single event carries little information and should be associated with low learning rates. In volatile environments, in contrast, a surprising outcome is more likely to signal a true change in the environment, and a single event is thus highly informative about current probabilities. Therefore, in changeable environments, individuals should rapidly update their beliefs about likely action-outcome contingencies instead of relying on the full history of events. Confirming this relationship between environment volatility and learning rates, it has been found that people's speed of learning tends to be higher in volatile than in stable conditions (Behrens et al., 2007).

2.2.2. The Information Bias Learning Task (IBLT)

Recently, it has been demonstrated that healthy adults can adjust their learning rates in response to volatility changes independently for positive and negative events (Pulcu & Browning, 2017). In this study, participants were asked to complete the computerised Information Bias Learning Task (IBLT), in which the volatility of associations between visual stimuli and positive (win) or negative (loss) outcomes was manipulated across the task blocks. It was found that healthy participants could simultaneously maintain high

learning rates for wins and losses if both outcomes were volatile. In addition, participants were able to uphold high learning rates for volatile wins concurrent with low learning rates for stable losses, and vice versa. As learning rates in this task have been proposed to reflect underlying beliefs about how informative positive and negative events are, these findings suggest that healthy participants can accurately estimate and utilise the information content of both positive and negative events independently to inform their decisions.

A first application of this task might be to identify potential basic information processing biases or impairments in psychiatric disorders. For example, prior research has indicated that people with high anxiety tend to adjust their learning rates to a smaller extent in response to volatility changes of the environment compared with non-anxious individuals (Browning et al., 2015). Furthermore, the IBLT could serve as an interventional paradigm by manipulating how informative positive and negative events are, and thereby teach individuals to attend to outcomes of a particular valence. Alternatively, this task may be used to track changes in reinforcement learning and affective biases following interventions. In the neurostimulation studies reported in this thesis, I examined whether it is possible to change people's learning rates for informative and uninformative positive or negative outcomes on the IBLT by altering activation in relevant neural networks with transcranial direct current stimulation (tDCS). In the next section, an overview of the practical application and physiological principles underlying the effects of this non-invasive brain stimulation technique is provided.

2.3. Physiological mechanisms of Transcranial Direct Current Stimulation (tDCS)

2.3.1. General overview

Non-invasive brain stimulation (NIBS) techniques such as transcranial direct current stimulation and transcranial magnetic stimulation aim to alter cortical excitability in targeted neural regions. In early research studies, NIBS was primarily used as a tool to

investigate potential causal contributions of specific brain structures to cognitive and motor functioning (Siebner & Rothwell, 2003). More recently, these methods are being investigated as potential treatment interventions for neurological and psychiatric conditions by modulation of abnormal neural activation, such as in depression (Boggio et al., 2008; Brunoni et al., 2011a; Dell'Osso et al., 2012; Fregni et al., 2006a; Loo et al., 2012), Parkinson's disease (Boggio et al., 2006; Fregni et al., 2006b; Wu et al., 2008) and post-stroke motor impairments (Hesse et al., 2011; Kim et al., 2010; Wu et al., 2013). A main advantage of NIBS versus other treatments such as pharmacotherapy is the capacity to change brain activation in anatomically specific regions (Bikson et al., 2012). Further clinical benefits of NIBS, and tDCS in particular, include its ease of use, safety and tolerability, and portability (Brunoni et al., 2012).

2.3.2. Administration of tDCS

In tDCS a weak direct electrical current (typically 1-2 mA) is delivered by a battery-powered device through two or more electrodes placed on the scalp and fixed to the head with elastic bands (Priori et al., 1998; Nitsche & Paulus, 2000). The electrodes are typically large (25-35 cm^2) rubber sheets encased in sponges (Ruffini et al., 2013) soaked in saline or covered in gel or conductive cream to reduce electrical resistance (see Figure 2.1.A.). One electrode(s) is placed over the target site of stimulation, with the other electrode(s) being placed on the scalp or an extracephalic region. The current flows from the anodal to the cathodal electrode(s), thereby generating weak electric fields (EFs; <1 V m^1) across the cortex (Datta et al., 2009; Salvador et al., 2010).

Electrodes can be positioned based on the international 10-20 electrode placement system as shown in Figure 2.1.B. (Jasper, 1958; Chatrian et al., 1985; Oostenveld & Praamstra, 2001). Alternatively, magnetic resonance imaging (MRI)-guided neuronavigation can be used to establish the target tDCS region (Dmochowski et al., 2013;

Fitzgerald et al., 2009). As the correspondence between skull landmarks and brain regions varies across individuals (Wörsching et al., 2016), MRI-guided localization is likely to improve accuracy of electrode positioning. While this method accounts for anatomical variability across individuals, however, financial and practical costs impede its widespread use in tDCS research.

Participant blinding is routinely achieved by either switching off the stimulation after an initial ramp-up period, or by using a sham stimulation mode on the tDCS device (Woods et al., 2016). To maintain blinding of the researcher, some tDCS devices include a study mode which displays realistic impedance values during sham stimulation and gives no indication of the true tDCS condition. Using these procedures, both participants and investigators are unable to distinguish between active and sham tDCS at 1 mA (Gandiga et al., 2006), although it has been suggested that blinding may be less effective at 2 mA due to more pronounced physical sensations and skin redness with higher current intensities (O'Connell et al., 2012).

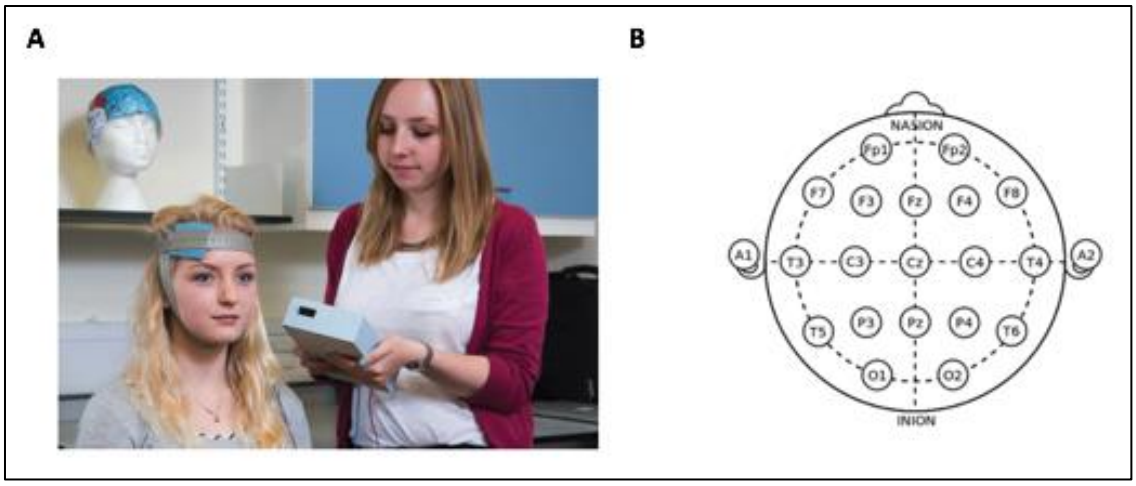

Figure 2.1. (A) Example of the tDCS equipment and application. The anodal and cathodal electrodes are inserted in sponges and attached to the scalp using rubber bands. The electrical current is delivered by a battery-operated device controlled by the research. **(B)** Electroencephalogram (EEG) 10/20 coordinates utilised to determine positioning of the electrodes.

2.3.3. Physiological mechanisms of action

2.3.3.1. Measurement of neuronal excitability

Effects of tDCS on neuronal excitation have been studied at cellular, neurochemical, and electrophysiological levels. Early animal studies demonstrated that anodal stimulation of the motor cortex increased evoked and spontaneous activity, whereas cathodal stimulation resulted in the opposite outcome (Bindman, Lippold, & Redfearn, 1962; Creutzfeldt, Fromm, & Kapp, 1962; Purpura & McMurtry, 1965). More recently, effects of tDCS on neuronal excitability have also been examined in humans (Antal et al., 2004; Antonenko et al., 2017; Kim et al., 2014; Liebetanz et al., 2002; Nitsche & Paulus, 2000; Nitsche et al., 2003; Stagg et al., 2009). Most work has focused on stimulation of the primary motor cortex (M1), at least in part due to the possibility of measuring the neurophysiological impact of tDCS on this region via electromyographic recordings of motor cortico-spinal tract excitability using TMS. When applied to the motor cortex, TMS can induce motor-evoked potentials (MEPs) in contralateral peripheral muscles. By measuring the MEP amplitude, stimulation threshold at which MEPs are elicited (motor threshold; MT), and input-output curve (I-O curve) slopes before, during, and after tDCS, alterations in the excitability of the underlying motor cortical tissue can be quantified. In contrast to eloquent motor areas, changes in excitability in the cognitive domain cannot be measured through myogenic responses (Jacobson, Koslowsky, & Lavidor, 2012) and are instead mainly assessed with neuroimaging methods or indirectly through cognitive task performance. As there have been relatively few neuroimaging studies on the physiological effects of tDCS in prefrontal regions, most of the research discussed here will therefore concentrate on the mechanisms of action during and after tDCS in motor regions. Although it is presently unclear to what extent these findings from M1 extrapolate to other regions,

it has been suggested that the main underlying neurophysiological processes are likely to be similar (Stagg & Nitsche, 2011).

2.3.3.2. Intra-stimulation effects

In general, anodal tDCS of the motor cortex increases the MEP amplitude, while cathodal tDCS decreases it (see Figure 2.2.A; Nitsche & Paulus, 2000). In addition, it has been demonstrated that the slope of the I-O curve, which provides an index of the excitability of larger neuronal populations, was increased with anodal stimulation and decreased with cathodal stimulation (Figure 2.2.B.; Nitsche et al., 2005). In contrast, the motor threshold was unchanged by either anodal or cathodal tDCS (Nitsche et al., 2005). Since motor threshold indexes pyramidal neuron excitability, this suggests that tDCS might affect interneurons more than pyramidal neurons.

At a cellular level, the dominant view is that these intra-stimulation effects depend on induced changes in polarisation of the resting membrane potential (Bikson et al., 2004; Kronberg et al., 2017; Nitsche & Paulus, 2000; Nitsche et al., 2003; Rahman et al., 2013). Evidence for this hypothesis has been provided by a study investigating the effects of pharmacological manipulations of voltage-dependent sodium and calcium channels on MEPs. Oral intake of the sodium channel blocker carbamazepine abolished the MEP amplitude increase normally observed during anodal tDCS, whereas the MEP decrease with cathodal tDCS was unchanged (Liebetanz et al., 2002; Nitsche et al., 2003). Similarly, administration of the calcium blocker flunarizine significantly diminished effects of anodal but not cathodal tDCS (Nitsche et al., 2003). These findings suggest that increases in excitability with anodal tDCS are caused by depolarisation of the membrane. In contrast, the absence of changes in the effects of cathodal tDCS are thought to reflect induced membrane hyperpolarisation via cathodal stimulation, thereby inactivating the calcium and sodium channels and consequently negating effects of carbamazepine and flunarizine.

Taken together, the evidence thus suggests that the intra-stimulation effects of anodal and cathodal tDCS are caused by membrane depolarisation or hyperpolarisation.

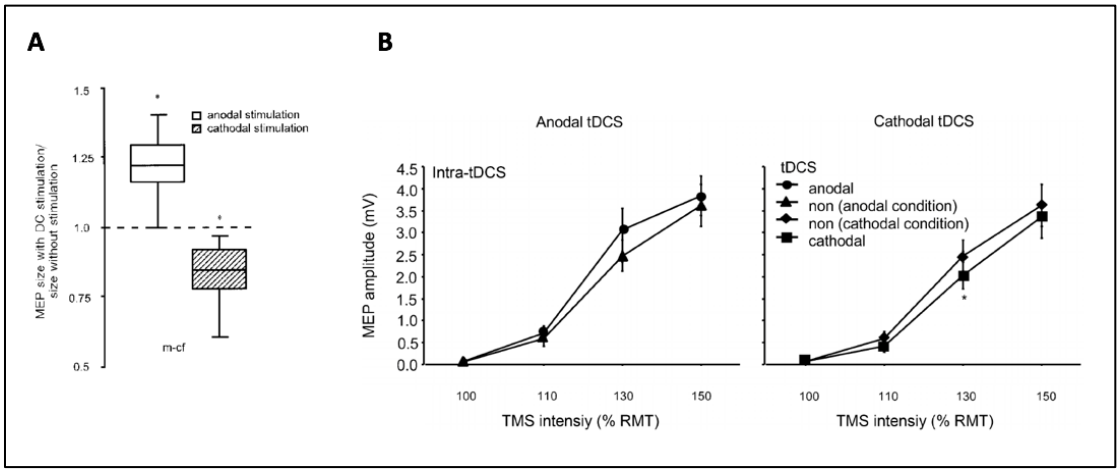

Figure 2.2. (A) Effects of anodal and cathodal tDCS at 1 mA on the MEP amplitude during stimulation (from Nitsche & Paulus, 2000). **(B)** MEP amplitudes (mV) at 100, 110, 130, and 150% of resting motor threshold (MT) during tDCS. Anodal tDCS at 2 mA increased the slope of the input-output curve, whereas cathodal tDCS decreased the slope (adapted from Nitsche et al., 2005).

2.3.3.3. After-effects of tDCS

The changes in evoked and spontaneous neural activity during tDCS described above can outlast the stimulation duration for several hours (Figure 2.3.A.; Bindman, Lippold, & Redfearn, 1962; Fritsch et al., 2010). This is of particular interest for tDCS in therapeutic contexts, where long-lasting induced changes in behaviour or cognition are desirable for treatment outcomes. The main mechanisms which have been proposed to drive the after-effects of tDCS are long-term potentiation (LTP) and long-term depression (LTD), which are both highly dependent on the activation of *N*-methyl-D-aspartate (NMDA) receptors that affect post-synaptic calcium influx (Bennett, 2000; Lüscher & Malenka 2012) and are thought to mediate various learning and memory processes (Aiba et al., 1994; Minichiello, 2009; Morris et al., 1986; Moser et al., 1998; Rioult-Pedotti, Friedman, & Donoghue, 2000; Ziemann et al., 2004). In mouse M1 brain slices, it was shown that pre-treatment

with an NMDA receptor antagonist D-APV (50 μM) prevented potentiation of field excitatory postsynaptic potentials (fEPSPs) normally elicited by direct current stimulation (Fritsch et al., 2010). In humans, it was similarly found that administration of the NMDA-receptor antagonist dextromethorphan completely eliminated post-stimulation changes in MEP amplitude after both anodal and cathodal tDCS (Figure 2.3.B.; Liebetanz et al., 2002; Nitsche et al., 2003).

In addition to NMDA activation, it has been suggested that brain-derived neurotrophic factor (BDNF) and activation of the receptor tyrosine kinase (TrkB) play a key role in the persistent excitability after-effects of tDCS (Fritsch et al., 2010; Minichiello, 2009). LTP has been found to increase release of BDNF, which in turn is a ligand for TrkB receptors that activate various intracellular signalling cascades (Minichiello et al., 2009). BDNF and consequent TrkB signalling have been shown to be vital to the after-effects of direct current stimulation. For example, in mice, both knockout of the BDNF Val66Met polymorphism and blockage of TrkB receptors were found to disrupt the augmentation effects of anodal tDCS on synaptic potentiation and motor skill acquisition (Figure 2.3.C.; Fritsch et al., 2010). Overall, the findings to date thus indicate that LTP- and LTD-like processes mediate the after-effects observed with tDCS.

Finally, evidence from magnetic resonance spectroscopy (MRS) studies suggests that the after-effects of anodal and cathodal tDCS may be associated with different neurochemical changes, although findings to date have been mixed. Anodal tDCS has been found to reduce GABA concentration in motor cortex, with no effect on Glutamate (Figure 2.3.D; Antonenko et al., 2017; Bachtiar et al., 2015; Kim et al., 2014; Stagg et al., 2009) although two recent studies did not replicate this finding (O'Shea et al., 2017; Petitet et al.; 2018). Cathodal stimulation has been shown to reduce both GABA and Glutamate in

one investigation (Stagg et al., 2009), but this result was again not replicated in other research studies (Antonenko et al., 2017; Kim et al., 2014).

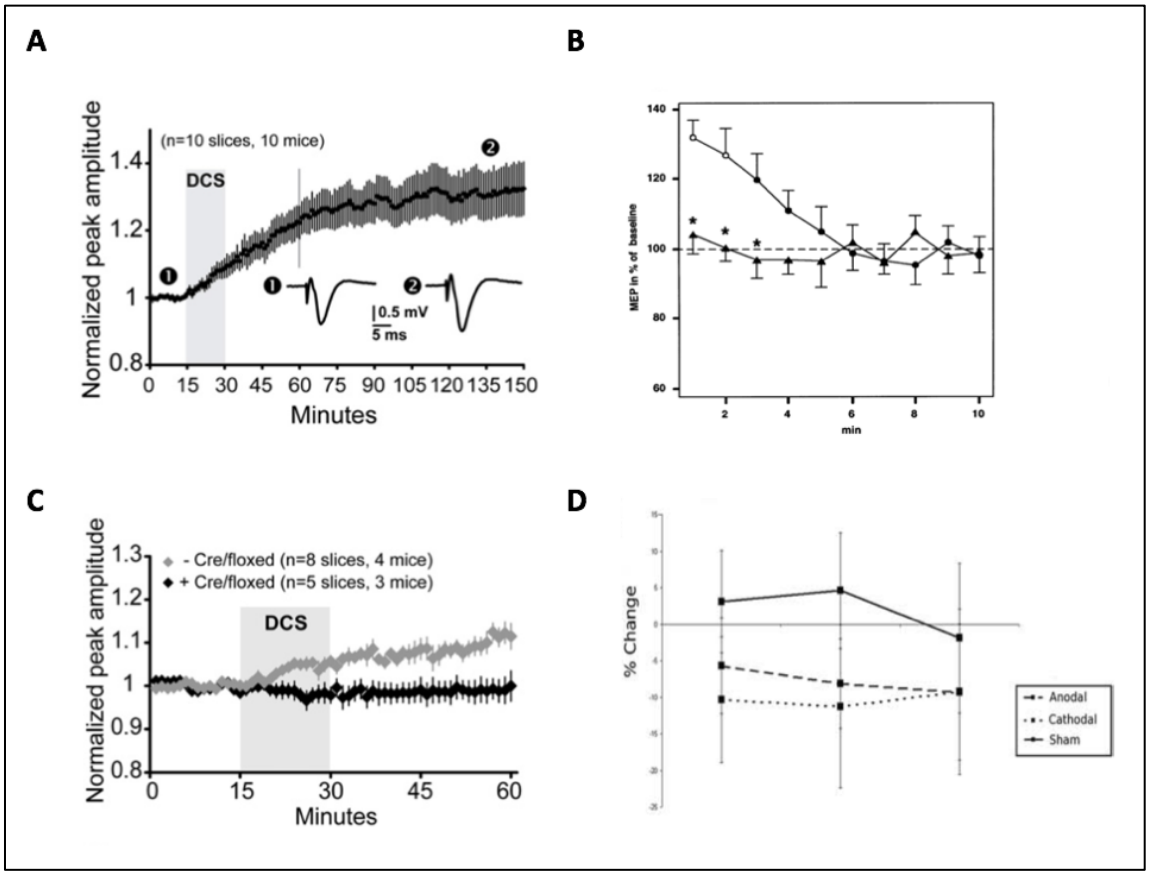

Figure 2.3. (A) Time course over two hours of field excitatory postsynaptic potential (fEPSP) amplitude following tDCS (Fritsch et al., 2010). **(B)** Post-stimulation MEP amplitudes after intake of dextromethorphan (triangles) or placebo (circles) as described by Liebetanz et al. (2002). **(C)** Absence of after-effects of direct current stimulation in mice with conditional BDNF knockout (black) and control type mice (grey; Fritsch et al., 2010). **(D)** Sustained increase in GABA concentrations up to twenty minutes after anodal tDCS of the primary motor cortex (Stagg et al., 2009).

2.3.3.4. Polarity-specific effects of tDCS

Whilst anodal stimulation in general increases cortical excitability and cathodal stimulation decreases it, the precise effects vary with specific parameter choices, including the duration, location, and intensity of tDCS. For example, whereas cathodal tDCS of the motor cortex at 1 mA for 20 minutes produced cortical inhibition, 2 mA stimulation

increased excitability (Batsikadze et al., 2013). In addition, Monte-Silva and colleagues (2013) showed that doubling the stimulation duration from 13 to 26 minutes at 1 mA shifted the effects of tDCS from excitation to inhibition. These findings indicate that increased intensity or prolongation of tDCS can produce non-linear effects. The authors suggest that this may be due to homeostatic mechanisms, such that neuronal counter-regulatory processes are activated in order to prevent potentially harmful overexcitation.

2.3.4. Spatial current distribution

In addition to the current intensity and duration, the influence of electrode placement on current distribution must be carefully considered when designing a tDCS protocol in order to optimise stimulation of the target regions. In general, the spatial resolution of tDCS is relatively diffuse (Bikson et al., 2012; Datta et al., 2009). The electrical field (EF) tends to decrease with distance from the electrodes, with peak activity occurring on the crowns of gyri and the bottom of sulci close to the electrodes (Figure 2.4.A.; Antal et al., 2017). Of significance, it has been reported that the folding of the cortex can produce inversions of tDCS polarity, such that both anodal and cathodal effects can be present in single gyri and sulci (Figure 2.4.B.; Datta et al., 2009). These models demonstrate that the electric field is not uniform, and alterations in behaviour with tDCS are therefore not likely to reflect homogenous excitability changes across the stimulated region. In addition, current flow can be affected by variability across individuals in brain and head anatomy factors such as skull thickness and cortical folding (Datta et al., 2009; Katz et al., 2007; Opitz et al., 2014; Russell et al., 2013; Wörsching et al., 2016).

To account for these fluctuations in the EF and the impact of individual anatomy, there has been a move towards the use of models to simulate the electric field induced by specific electrode montages. Computational forward models can range in complexity from concentric sphere models to high-resolution models based on individual's MRI scans

(Bikson et al., 2013). The latter can account for individual distortions in the current flow, and thereby assist in customising tDCS montages. Among others, these models have suggested that focality of tDCS can be enhanced with the use of high-definition tDCS, in which a small anodal electrode is positioned over the target region surrounded by four cathodal electrodes (Datta et al., 2009; Edwards et al., 2013; Minhas et al., 2010). In future studies, such computational simulations of the electric current may be used to optimise conventional tDCS approaches (Bikson et al., 2013).

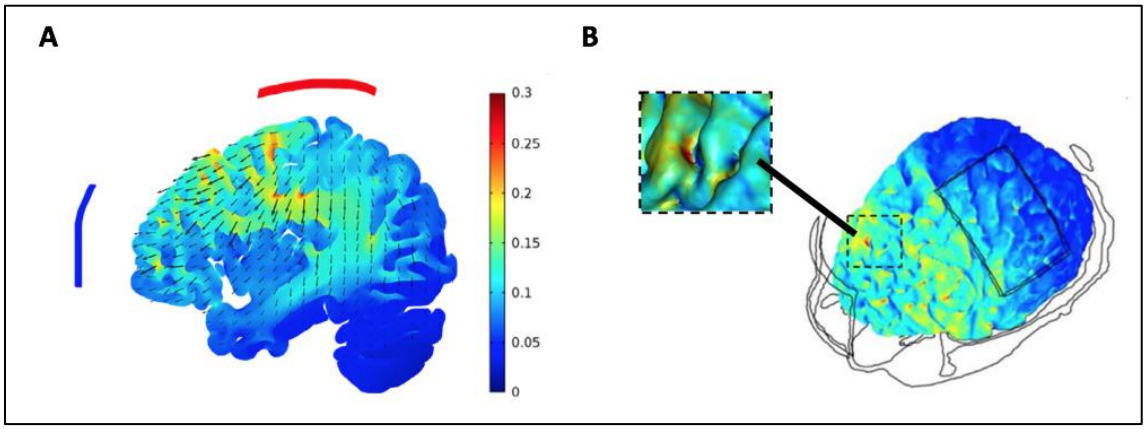

Figure 2.4. (A) Magnitude of the electric field in the cortex in V/m with anodal tDCS of the primary motor cortex, indicating hotspots on the peaks of the gyri and bottom of sulci (from Antal et al., 2017) **(B)** Magnitude plot of the electric field along the brain surface during tDCS with the anode over the left motor cortex and the cathode over the right contralateral orbita. As shown in the expanded image, anodal and cathodal stimulation can occur within gyri and sulci (adapted from Datta et al., 2009).

2.3.5. Safety considerations

The safety and tolerability of tDCS have been investigated in both healthy adults and psychiatric populations (Antal et al., 2017; Bikson et al., 2016; Poreisz et al., 2007). The most commonly reported side effects of tDCS include tingling and itching sensations under the electrodes, headache, and tiredness (Poreisz et al., 2007; Fertonani, Ferrari, & Miniussi, 2015). In terms of serious adverse events, an examination of over 33,000 sessions in more than 1,000 participants found no evidence for irreversible injury with

tDCS protocols applying a wide range of stimulation parameters (Bikson et al., 2016). In addition, tDCS has been found to offer a good tolerability profile in depressed patients (Dell'Osso et al., 2012; Meron et al., 2015). In the field of depression research, a potential adverse effect which may occur with prefrontal stimulation is instances of (hypo)mania. Individuals undergoing this tDCS protocol should therefore be carefully assessed for a history of bipolar disorder or manic episodes, as these factors may signal a higher risk of inducing a manic switch with tDCS (Antal et al., 2017). Overall, tDCS protocols using well-defined parameters seem to be safe and well-tolerated, and may present a novel opportunity for interventions of disorders characterised by aberrant neuronal activity.

2.4. Rationale for the tDCS protocol

2.4.1. Neural target for stimulation

Well-informed parameter choices are thus key to the design of studies employing transcranial direct current stimulation. First of all, an appropriate target region for stimulation must be identified. In terms of learning rates, a potential candidate for stimulation is the dorsal anterior cingulate cortex (dACC). Previously, it has been demonstrated that volatility estimates and associated learning rates for rewarding outcomes are linked to activity in a circumscribed region in the dACC (see Figure 2.5.A.; Behrens et al., 2007; Behrens et al., 2008). In addition, it has been proposed that the ACC is critical for encoding rewards and using reinforcement histories to guide behaviour (see Akitsuki et al., 2003; Ernst et al., 2004; O'Reilly et al., 2013; Rushworth et al., 2007; Wu et al., 2017). Supporting this hypothesis, it has been found that inhibition of ACC activity disables reward history-guided behaviour in both rats and macaque monkeys (Amiez, Joseph, & Procyk, 2006; Kennerley et al., 2006; Tervo et al., 2014). Humans with strong neural representations of the information to be learnt in the dACC, on the other hand, have been found to be better at using learnt information to guide their choices (Scholl et al.,

2017). Similarly, greater activation of the dACC has been associated with enhanced reward learning on a probabilistic learning task (Santesso et al., 2008). Crucially, depressed individuals have presented with dysfunctional activity in the dACC (Davidson et al., 2002; Mayberg, 2003), which may contribute to impairments in integrating reinforcement history over time.

The dACC therefore presents a suitable target for neurostimulation when attempting to change reward learning behaviour. Although the current induced by tDCS tends to affect superficial neural regions rather than stimulating deep brain structures directly, it has been shown that non-invasive brain stimulation can indirectly affect distal regions which are functionally or anatomically connected to the stimulated area (Chib et al., 2013; Hone-Blanchet, Edden, & Fecteau, 2016; Wörsching et al., 2016; Zheng, Aslop, & Schlaugg, 2011). It has previously been demonstrated that the lateral prefrontal cortex is strongly connected to the dACC (see Figure 2.5.B.; Beckmann, Johansen-Berg, & Rushworth, 2009; Pandya & Yeterian, 1991; Petrides & Pandya, 1999; Petrides & Pandya, 2002), and it may therefore be possible to alter dACC activation through stimulating frontal regions. Indeed, prior research has shown that repetitive transcranial magnetic stimulation (rTMS) of the mid-dorsolateral frontal cortex altered activity in the ACC (Paus, Castro-Alamancos, & Petrides, 2001). In addition, Weber et al. (2014) reported that anodal tDCS of the right DLPFC increased ACC activity as shown with magnetic resonance imaging scans. Similarly, an electroencephalogram (EEG) study indicated that tDCS over the left DLPFC was associated with increased amplitudes of the N450 component, which is thought to be generated by activity in the ACC, during an emotional cognitive control task (Vanderhasselt et al., 2013).

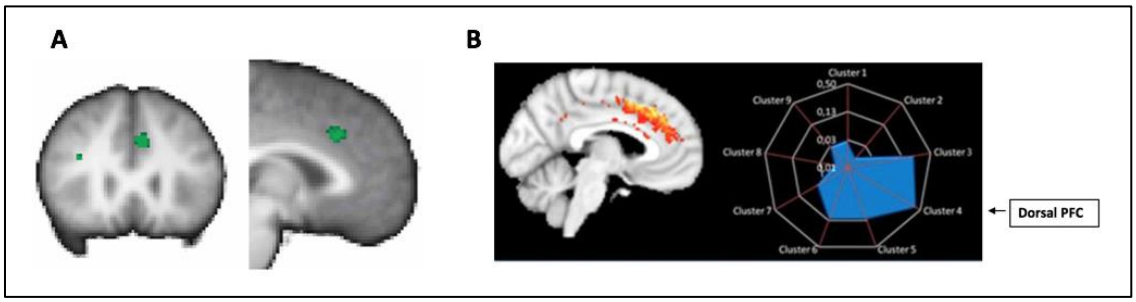

Figure 2.5. (A) dACC region correlated with estimates of the volatility of the environment, reproduced from Behrens et al. (2007). **(B)** Connectivity fingerprint of the anterior cingulate cortex, indicating that the dACC is most strongly connected with the dorsolateral prefrontal cortex (Cluster 4; from Beckmann et al., 2009).

2.4.2. Stimulation parameters

In the neurostimulation studies reported in this thesis, a bifrontal tDCS protocol commonly employed in depression research was utilised. In this montage, with the anode over the left DLPFC and the cathode over the right DLPFC, tDCS is applied at 2 mA for 20 minutes. These electrode placement and current intensity parameters have been used successfully to reduce symptoms of depression in a number of tDCS studies (Brunoni et al., 2011a; Brunoni et al., 2013a; Brunoni et al., 2013b; Dell'Osso et al., 2012; Dell'Osso et al., 2014; Ferrucci et al., 2009; Valiengo et al., 2013; but see Blumberger et al., 2012 for a report of non-significant changes in symptoms of severely depressed patients). While there is thus evidence to suggest that this procedure is effective in moderating general symptoms of depression, it has not yet been tested whether bifrontal tDCS can alter specific reward learning processes. Modelling studies simulating direct electric current distribution in healthy adults and MDD patients have indicated that the electric field induced with bifrontal tDCS is capable of reaching the ACC (Bai et al., 2014; Csifcsák et al., 2018). It was therefore expected that tDCS of bilateral DLPFC might be capable of altering learning rates through indirect modulation of the dorsal anterior cingulate cortex. This hypothesis was tested in a series of experimental tDCS studies described in Chapters 4-6.

Chapter 3. Examining test-retest reliability and practice effects on the Information Bias Learning Task (IBLT)

3.1. Introduction

In the Information Bias Learning Task (IBLT), a reinforcement learning paradigm, the information content of win and loss outcomes is manipulated by varying the volatility of these outcomes independently (see Section 2.2.2.; Behrens et al., 2007; Browning et al., 2015; Pulcu & Browning, 2017). By quantifying how quickly individuals learn from informative and uninformative positive or negative outcomes, respectively, this task can provide a measure of affective biases in learning processes. Using various versions of the IBLT, it has been demonstrated that healthy adults preferentially learn from the most informative outcomes (Pulcu & Browning, 2017), whereas this tendency may be reduced in individuals with high anxiety (Browning et al., 2015).

The IBLT has thus been applied to investigate reinforcement learning mechanisms across individuals, as well as to assess differences in learning between groups based on symptoms of psychiatric disorders. At present, however, it is unclear how stable performance on this task is within individuals across testing sessions. Changes in task scores over time can not only be caused by true alterations in the measured ability or trait, but also by factors such as measurement errors and practice effects (Harvill, 1991). Knowledge about the consistency of IBLT performance over time is therefore crucial for correct interpretations of findings from repeated-measures studies. Moreover, a better understanding of the psychometric properties of this task, such as test-retest reliability, may clarify its potential as a diagnostic and/or prognostic tool.

Test-retest reliability represents the stability of task performance and the underlying construct being measured, as well as the influence of errors on the measurement over multiple sessions. There can be various sources of error contributing to

the overall score, including simple mistakes, fatigue, or lucky choices (Weir, 2005). If the error across measurements is greater than actual changes in the score, interpreting findings from longitudinal studies becomes challenging. It is therefore important to have an estimate of the influence of error on observed outcomes in cognitive tasks. As true scores and measurement errors cannot be differentiated directly, they are generally estimated using statistical measures. In psychological research, reliability is often measured using intraclass correlation coefficients (ICCs) and standard errors of measurement (SEMs; Weir, 2005), which reflect the correlation and agreement between multiple measurements. Practice effects are a specific source of systematic error, in which performance consistently changes across measurements due to factors such as enhanced strategies, increased familiarity with the test-taking environment, or improved understanding of the task (Lo et al., 2012). In addition to overall test-retest reliability, an awareness of potential practice effects is therefore essential to ensure that enhanced performance is not misconstrued as an intervention effect or clinical improvement in clinical trials.

In the present study, I aim to characterise the test-retest reliability of and practice effects on the IBLT across two sessions in healthy adults. First, I will examine whether the previous finding that healthy adults are able to adjust their estimates of the information content for positive and negative events independently on the IBLT as indicated by learning rates is replicated (Pulcu & Browning, 2017). It was expected that participants would have high learning rates for volatile outcomes, whereas learning rates would be lower for stable outcomes across the task. Furthermore, I endeavoured to characterise test-retest reliability on the IBLT as indicated by intraclass correlation coefficients and standard errors of measurement. Based on previous guidelines, ICC values were classified as poor (below 0.5), moderate (between 0.5 and 0.75), good (between 0.75 and 0.9), or excellent (greater than 0.9) test-retest reliability (Koo & Li, 2016). Finally, I examined

whether participants showed practice effects on the IBLT across two testing sessions using analysis of variance models.

3.2. Methods and Materials

3.2.1. Participants

Thirty healthy participants between the ages of 18 and 45 years (21 female) were recruited from the community via advertisements and word of mouth. Ethical approval was obtained from the Central University Research Ethics Committee (CUREC R48995/RE002) and written informed consent was given by all participants at the start of the first testing session. Participants were reimbursed for travel expenses and time at a rate of £10 per hour. In addition, all participants received the total amount won in the computerised task performed during the testing sessions.

3.2.2. General procedure

Participants were invited to two testing sessions, which took place at least one week apart. In the first session, participants were asked to fill out Beck's Depression Inventory II (BDI; Beck, Steer, & Brown, 1996) and the Trait version of the State-Trait Anxiety Inventory (STAI-Trait; Spielberger et al., 1983). The score ranges were 0 – 63 on the BDI and 20 – 80 on the STAI-Trait, with higher scores indicating lower mood and higher anxiety, respectively. Next, the researcher explained the computerised task (described below) called the Information Bias Learning Task or IBLT, which was presented on a laptop monitor using *Presentation* software version 18.3 (Neurobehavioral Systems Inc, Albany, CA). Participants were given the opportunity to practice the task and ask further questions if needed. The participant was then given a pair of headphones and was asked to complete 6 task blocks of the IBLT.

3.2.3. Information Bias Learning Task (IBLT)

3.2.3.1. Overview

The Information Bias Learning Task was adapted from learning tasks previously used to investigate the effects of volatility on reward and loss learning rates (Behrens et al. 2007; Browning et al., 2015; Pulcu & Browning, 2017). In the IBLT, participants are instructed to try and win as much money as possible by choosing between two visual stimuli which are probabilistically associated with a win and/or a loss outcome. The occurrence of the win and loss outcomes are independent, such that the association of a win outcome with a particular shape does not determine the association of the loss outcome and vice versa. Predicted shape-outcome associations are learned over time through experience on previous trials. In each task block, the volatility of choice-outcome contingencies of win and loss outcomes are modulated independently. In a volatile state, the association of stimulus A and a win (or loss) outcome switches between 20% and 80% in runs of 14-30 trials. In stable conditions, on the other hand, the association between stimulus A and a win (or loss) outcome remains constant at 50%. In "Win-volatile" blocks, win outcomes are volatile while loss outcomes are stable, whereas in "Loss-volatile" blocks loss outcomes are volatile and win outcomes are stable. In "Both-volatile" blocks, both the win and loss outcome associations are volatile.

3.2.3.2. Stimuli

The IBLT was presented on a Dell Latitude E5510 laptop and consisted of six blocks of 80 trials each, with a brief resting session of 30 seconds between each of the blocks. On every trial, participants were presented with a fixation cross in the centre of the screen flanked by two abstract shapes (letters selected from the Agathodaimon font) offset by approximately 7° visual angle. Participants were instructed to choose one of the shapes through a button press on a keyboard. Following each choice, a win and loss outcome

would appear above or below the shapes (Figure 3.1.A.) in randomised order for an interval of 1000 ms. Win and loss outcomes for the chosen shape were accompanied by positive ("cha-ching") and negative (error buzz) auditory stimuli lasting 0.7 seconds. Win and loss outcomes were independently associated with the two shapes, such that a shape could be associated with one, both, or neither of the outcomes. Choice of a shape associated with a win could result in a gain of 10p, whereas choosing a shape associated with a loss outcome led to a loss of 10p. If the participants chose a shape that had both outcomes or neither a win nor a loss, the amount won remained unchanged. Participants' accumulated total winnings were displayed under the fixation cross and were updated at the beginning of the subsequent trial (see Figure 3.1.B.). At the start of the task, the total amount was fixed at £1.50 for each participant.

Within each of the six task blocks the same two shapes were used for all trials, with different shapes being used between the task blocks. Volatility of the win and loss outcomes was manipulated across the blocks (see Figure 3.1.C). The first and last blocks were "Both-volatile", with shape-outcome associations being volatile for both win and loss outcomes. Blocks 2 and 3 were "Win-volatile" and "Loss-volatile", respectively, with the order of the blocks being counterbalanced across participants. After the third block the task was re-started, with the task blocks being presented in a different order. Blocks 4 and 5 were "Win-volatile" and "Loss-volatile", and were presented in the same order as Blocks 2 and 3 (see Figure 3.2.). Behavioural responses on the task were analysed by fitting the computational models described in section 3.2.4.

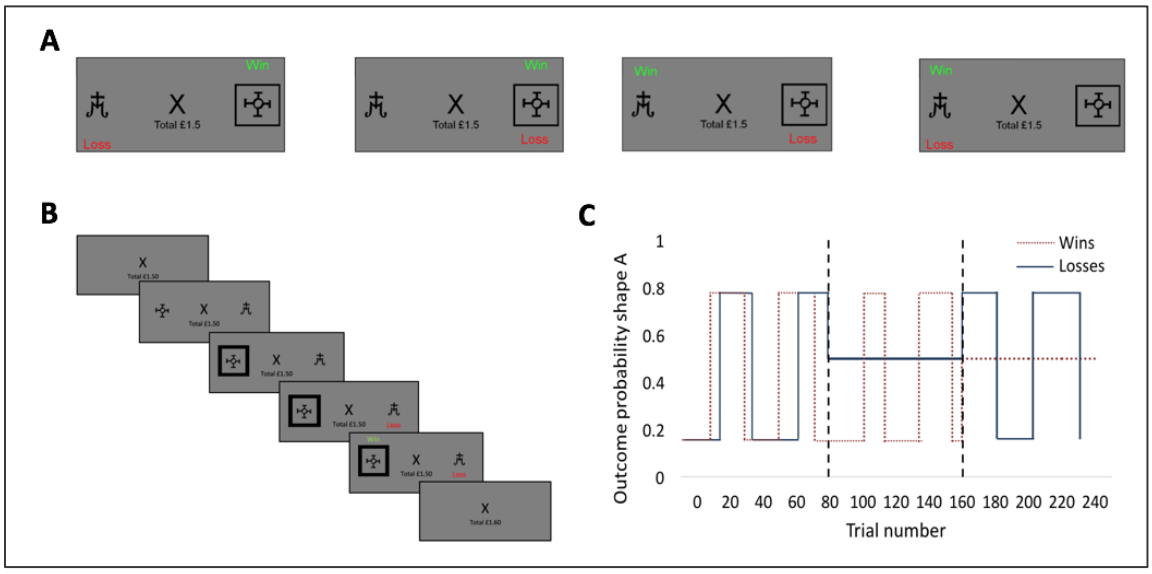

Figure 3.1. Schematic illustration of the Information Bias Learning Task. (A) Win and loss outcomes are manipulated independently, such that a shape can be associated with a win, a loss, both, or neither. **(B)** After showing the fixation cross and total amount of money won, two abstract shapes are presented on either side of the cross. Once the participant has chosen one of the shapes, a black frame appears around that shape and a win and loss outcome appear successively in randomised order. The total amount of money is updated at the start of the next trial. **(C)** Volatility of the win and loss outcomes is manipulated independently across task blocks, either switching between 20% and 80% choice-outcome associations or remaining stable at 50% choice-outcome associations.

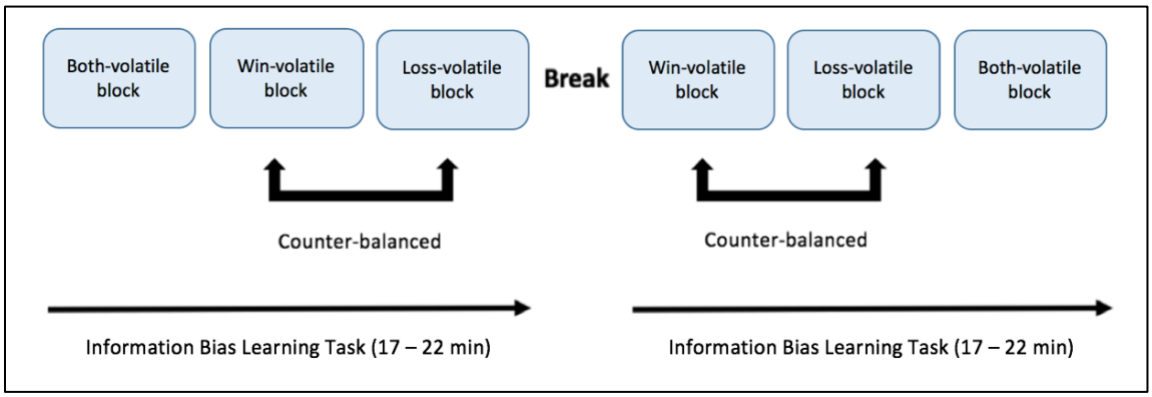

Figure 3.2. Order of the six IBLT task blocks performed by each participant, with outcome volatility being manipulated across the blocks.

3.2.4. Computational modelling and choice behaviour

The primary outcome of interest in the IBLT is the learning rate for win and loss outcomes across the six blocks. A behavioural model (Model 1), similar to that previously used by Pulcu and Browning (2017), was used to estimate learning rates for each participant using a commercial software package (MATLAB R2016B, The Mathworks Inc, Natick, MA).

Model 1: This model captures changes in learning rates associated with volatility, as well as the influence of choice stochasticity on behaviour as measured with an inverse temperature parameter. First, the model calculates the probability estimates that the win and loss outcomes would be associated with shape "A" using a Rescorla-Wagner learning Rule (Rescorla & Wagner, 1972):

$$rwin_{(i+1)} = rwin_{(i)} + \alpha win * \varepsilon_{win(i)} \quad (3)$$

$$rloss_{(i+1)} = rloss_{(i)} + \alpha loss * \varepsilon_{loss(i)}$$

in which $rwin_{(i+1)}$ is the estimated win outcome probability for the i+1st trial, $rwin_{(i)}$ is the estimated outcome for the ith trial, αwin represents the learning rate for win outcomes, and $\varepsilon_{win(i)}$ indicates the prediction error on the ith trial. Prediction error is calculated as the predicted outcome value minus the actual outcome value. At the start of each block, *rwin* was initialised at 0.5 as participants could not have prior expectations about which shape was most likely to be associated with a win outcome. $rloss_{(i+1)}$, $rloss_{(i)}$, *aloss* and $\varepsilon_{loss(i)}$ indicate the same variables for the loss outcome. Next, estimated outcome probabilities were transformed into a single choice probability using a softmax function:

$$P_{(choice=A(i))} = \frac{1}{1+\exp^{(-\beta*(rwin_{(i)}+t-rloss_{(i)}))}} \quad (4)$$

Here, $P_{(choice=A(i))}$ is the probability of choosing shape A in trial *i*, with β representing the inverse decision temperature and *t* reflecting an added parameter used to estimate a general tendency to select one of the options over the other. The inverse temperature indicates the degree to which the expected values are used to determine choice for a particular shape. Learning rates and β-values were calculated separately for each task block and participant. This was achieved by calculating the full joint posterior probability of the parameters given participants' choices, deriving the expected value of each parameter from their marginalised probability distributions. The first 10 trials of each block were omitted when fitting the model parameters to participants' choices, as initial learning rates without prior knowledge were expected to differ from informed learning rates in later trials of the task.

The choice for this particular model was based on a comparison of six different computational models for all experimental studies in this thesis, which showed that the present model provided the best fit to the data as indicated by Akaike Information Criterion (AIC) and Bayesian Information Criterion (BIC) values for the individual models. Specifically, AIC and BIC values were calculated for the current model and the following five alternative models (summarised in Table 3.1):

Model 2: Rather than modelling a single inverse temperature value, it is possible that participants display different stochastic choice behaviour for the win and loss outcomes. This model, which has been applied in previous studies of the IBLT (Pulcu & Browning, 2017), therefore incorporates separate inverse temperature parameters for wins and losses as follows:

$$P_{(choice=A(i))} = \frac{1}{1+\exp^{(-\beta win * rwin_{(i)} - \beta loss * rloss_{(i)})}} \quad (5)$$

Model 3: Instead of learning the independent probabilities of win and loss outcomes, participants might take a model-free approach to the task by learning an overall value of each of the presented shapes (Pulcu & Browning, 2017):

$$v^{A}{}_{(i+1)} = v^{A} + \alpha * (out_{(i)} - v^{A}_{(i)}) \quad (6)$$

in which v^{A} represents the value of shape "A", α is a single learning rate for updating the value, and $out_{(i)}$ is the outcome of trial *i* (i.e., win – loss for shape "A", which can be -1, 0, or 1). In the first trial, the value of shape "A" is set at 0. The estimated values of the two presented shapes were transformed into a choice probability by applying a softmax function using a single inverse temperature parameter.

Model 4: Similar to Model 1, this model calculates two learning rates and one inverse temperature parameter. This model is slightly simpler, however, in that it omits the "tendency" parameter *t* as described in Equation 4.

Model 5: A model using a single learning rate for win and loss outcomes combined, with two inverse temperature values for the two individual outcomes.

Model 6: The final model was similar to Model 2, with the main difference being that the values for the win and loss outcomes were centred at zero prior to multiplication with the inverse temperature values:

$$P_{(choice=A(i))} = \frac{1}{1+\exp^{(-(\beta win*(rwin_{(i)}-0.5))-(\beta loss*(rloss_{(i)}-0.5))}} \quad (7)$$

As shown in Figure 3.3., Model 1 provided the best fit to the data across the studies (note that lower AIC and BIC values are indicative of a better model fit). When investigating comparator models for the individual research studies in isolation, it was similarly found that Model 1 had the closest fit to the data in each of the experiments (see Appendix A.).

Table 3.1. Description of comparator models of the IBLT.

Model	Number of learning rate parameters	Number of inverse temperature parameters	Description
1	2	1	Model used in analyses
2	2	2	Two inverse temperature parameters model
3	1	1	Model-free learner
4	2	1	Single inverse temperature model omitting the tendency parameter
5	1	2	Single learning rate model
6	2	2	Reduced win and loss values

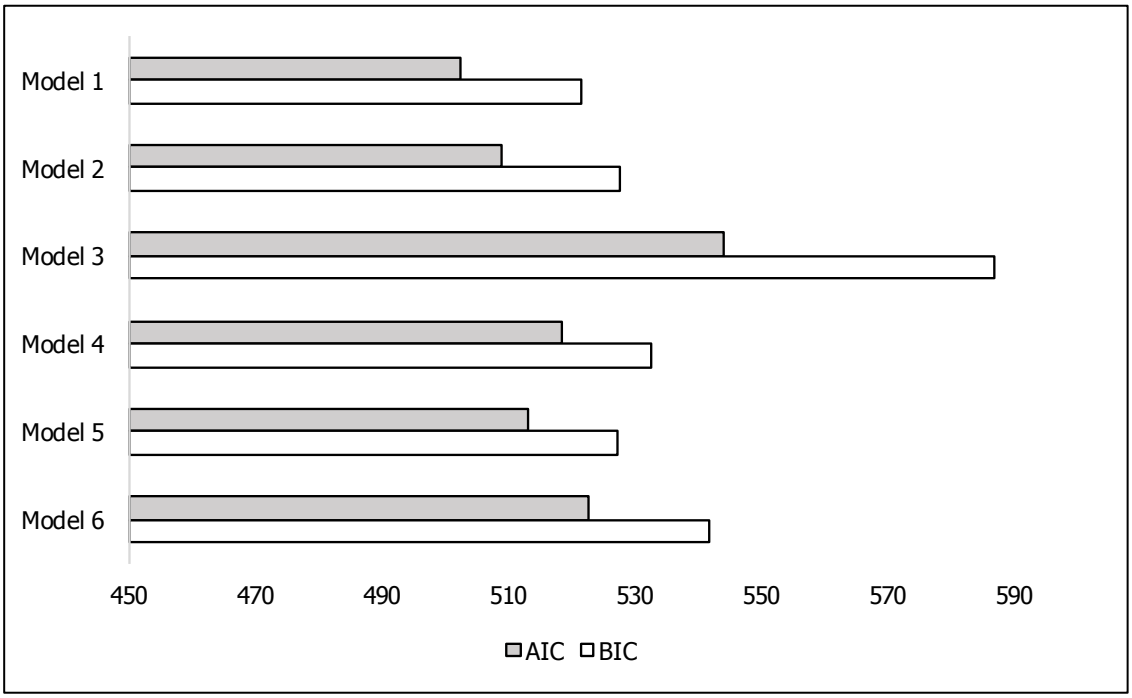

Figure 3.3. Average sum of the Akaike Information Criterion (AIC) and Bayesian Information Criterion (BIC) values for comparator models of the Information Bias Learning Task across the experimental studies (Chapters 3-6) per participant. Lower values indicate a better fit of the model to the data.

3.2.5. Win-driven choice behaviour

To complement the model-based analyses, non-computational analyses were conducted to assess choice behaviour for win and loss outcomes. Specifically, I examined the impact of win and loss outcomes on choices following trials where a shape was associated with both a win and a loss. In such trials, the win outcome encourages the participant to select the shape that had both a win and loss in trial *i*, whereas the loss outcome encourages the participant to select the shape that had neither a win nor a loss. Therefore, the shape selected by the participant on trial *i*+1 provides an indication of whether the positive or negative outcome has a stronger impact on subsequent choice behaviour (see Figure 3.4.). Win-driven choice behaviour was quantified as the sum of the following choices:

a) If the chosen shape on trial *i* had both a win and a loss, participants selected the same shape on trial *i*+1
b) If the chosen shape on trial *i* had neither a win nor a loss, participants selected the other shape on trial *i*+1

The proportion of win-driven choices was then calculated by dividing the number of win-driven choices by the total amount of trials on which a shape was associated with both a win and a loss (40 per block).

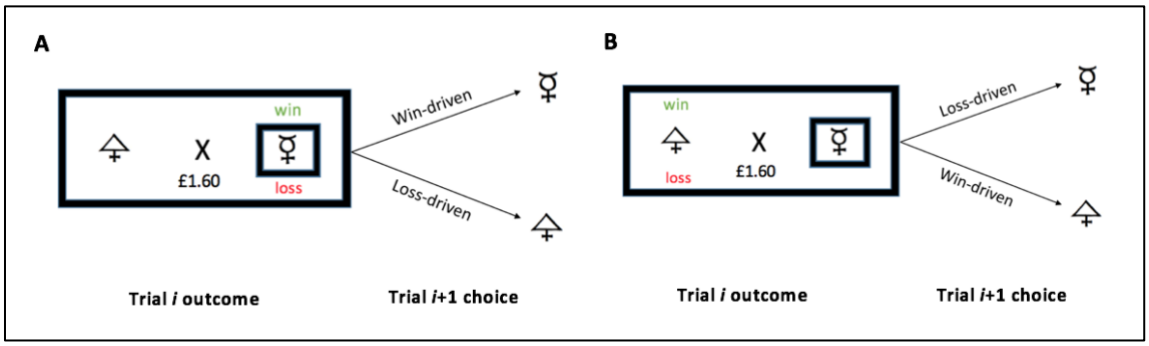

Figure 3.4. Schematic illustration of win-driven behaviour on the IBLT. **(A)** If the participant had chosen a shape associated with both a win and a loss on trial *i*, choosing the same shape on the next trial *i* + 1 indicated a stronger impact of win than loss outcomes on behaviour. **(B)** Conversely, if the participant had chosen a shape with neither a win nor a loss on trial *i*, choosing the alternative shape on the next trial *i* + 1 suggested a greater influence of win than loss outcomes on choice behaviour.

3.2.6. Statistical analyses

Statistical analyses were carried out using R software (R Version 3.3.1.). In order to assess test-retest reliability, intraclass correlation coefficients were computed using the {irr} package. A oneway model was applied, which calculated consistency ICCs while considering participants as random effects. First, ICCs of the learning rates were calculated separately for win and loss outcomes for each of the six blocks completed by the participants. Next, coefficients were calculated to determine the reliability of the volatility effect for wins and losses separately. These values were determined by subtracting the learning rates for stable outcomes (e.g. wins in the Loss-volatile block) from the learning rates for volatile outcomes (e.g. wins in the Win-volatile block). These coefficients provide information on whether learning from informative (volatile) compared to non-informative (stable) outcomes is consistent across a one-week interval. In other words, this measure indicates how reliably participant choices tend to be influenced by the information content of either win or loss outcomes. Finally, potential asymmetries in learning from win versus loss outcomes were investigated by subtracting loss learning rates from win learning rates

for the Win-volatile and Loss-volatile blocks separately. For the β-values and win-driven choice behaviour, ICCs were calculated for each of the six blocks across the two sessions. In addition, standard error of measurement (SEM) values were calculated for each of the blocks and outcomes using the {psychometric} package, which applies the following equation:

$$SEM = s * sqrt(1 - rxx) \quad (8)$$

in which *s* is the standard deviation in test scores on the IBLT and *rxx* represents the ICC values for the test-retest reliability. These SEM values represent the amount of variation in measurement errors on the IBLT, and can be used to calculate a confidence range around the observed test scores (Harvill, 1991).

Potential practice effects were examined by carrying out repeated-measures ANOVAs for the learning rates, inverse temperature values, and win-driven choice behaviour. For the learning rates, the within-subject predictor variables used were Block (Win-volatile, Loss-volatile, or Both-volatile), Outcome valence (win versus loss), Run (first 3 blocks versus second 3 blocks within a session), and Session (first versus second). Similarly, predictor variables for the inverse temperature values and win-driven choice behaviour were Block, Run, and Session. An inverse logit transform was applied to the learning rates while a log transform was used on inverse temperatures, so that the outcome values were on the infinite real line. Figures and reported values represent raw parameter values to facilitate interpretation of the results.

Finally, in the present study as well as the neurostimulation research studies (Chapters 4-6), I tested whether trait mood and anxiety scores were associated with win and loss learning rates at baseline to assess if these subjective measures were predictive of

information processing biases. As these analyses were not significant in any of the studies, however, the results are reported as supplementary data (see Appendix B) in the interest of brevity.

3.3. Results

3.3.1. Task performance

The first aim of the present study was to examine whether previous reports of adaptation of learning rates to volatility of the environment could be replicated. Repeated-measures ANOVAs of the learning rates indicated a main effect of Block ($F(2,58) = 9.34$, $p < 0.001$, $\eta^2 = 0.022$) and Outcome valence ($F(1,29) = 9.23$, $p < 0.001$, $\eta^2 = 0.016$). In addition, there was a Block x Outcome valence interaction ($F(2,58) = 27.05$, $p < 0.001$, $\eta^2 = 0.128$) as shown in Figure 3.5. Paired *t*-tests indicated that learning rates were higher for win outcomes than for loss outcomes in the Win-volatile blocks ($t(119) = 11.51$, $p < 0.001$). In addition, learning rates were higher for win outcomes in the Win-volatile than the Loss-volatile blocks ($t(119) = 7.11$, $p < 0.001$), while loss learning rates were higher in the Loss-volatile than the Win-volatile blocks ($t(119) = 8.95$, $p < 0.001$). Finally, win learning rates were higher in the Both-volatile blocks than the Loss-volatile blocks ($t(119) = 6.10$, $p < 0.001$), and loss learning rates in the Both-volatile blocks were greater than in the Win-volatile blocks ($t(119) = 6.44$, $p < 0.001$). These results confirm previous findings that healthy participants tend to have higher learning rates for volatile than stable outcomes, and can maintain separate learning rates for positive and negative events (Pulcu and Browning, 2017).

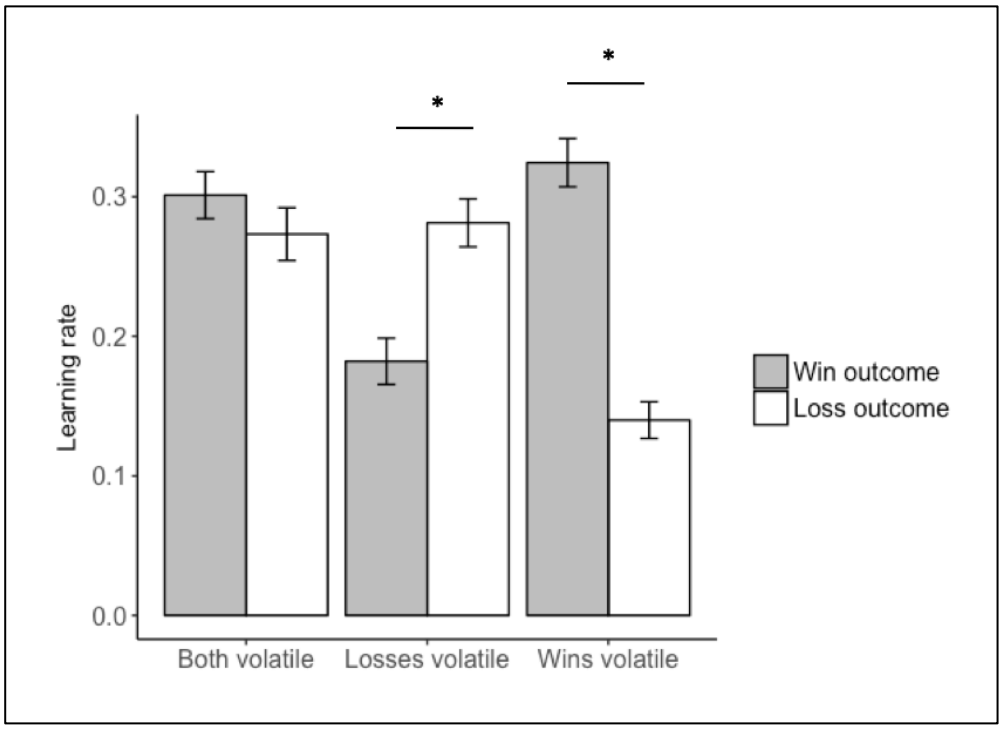

Figure 3.5. Effect of the volatility manipulation on mean learning rates on the IBLT. Pair-wise comparisons indicated that both win and loss learning rates tended to be higher if the outcomes were volatile than if they were stable (*p <0.001). Error bars represent the standard error of the mean.

3.3.2. Test-retest reliability

Next, test-retest reliability of the Information Bias Learning Task was investigated. ICC values, 95% confidence intervals (CI), and the SEM of the intraclass correlation coefficients were calculated for each of the task blocks for win and loss learning rates (Table 3.2), as well as the difference between win – loss learning rates (Table 3.3.) and the difference between learning rates for volatile – stable outcomes (see Table 3.4.). ICC values were also estimated for the inverse temperature values (Table 3.5.) and win-driven choice behaviour (Table 3.6.). As can be seen from these tables, ICC values ranged from poor to moderate across the different outcomes and task blocks.

Table 3.2. IBLT test-retest reliability estimates of the learning rates for win and loss outcomes.

Outcome	Block	Run	ICC	95% CI	SEM
Wins	Wins volatile	One	0.662	0.404 – 0.823	0.625
		Two	0.571	0.275 – 0.769	0.683
	Losses volatile	One	0.071	-0.287 – 0.414	1.227
		Two	0.186	-0.176 – 0.506	1.137
	Both volatile	One	-0.013	-0.363 – 0.342	1.109
		Two	-0.140	-0.468 – 0.225	1.664
Losses	Wins volatile	One	0.416	0.076 – 0.671	0.890
		Two	0.301	-0.057 – 0.591	0.895
	Losses volatile	One	0.385	0.039 – 0.650	1.025
		Two	0.348	-0.004 – 0.625	0.834
	Both volatile	One	0.374	0.026 – 0.642	1.065
		Two	0.631	0.360 – 0.805	0.947

ICC = Intra-class correlation coefficient; CI = Confidence interval; SEM = Standard error of measurement

Table 3.3. Test-retest reliability of win – loss learning rates on the IBLT.

Block	Run	ICC	95% CI	SEM
Wins volatile	One	0.718	0.491 – 0.855	0.762
	Two	0.190	-0.173 – 0.509	1.357
Losses volatile	One	0.355	0.004 – 0.629	0.048
	Two	0.437	0.102 – 0.685	0.039
Both volatile	One	-0.004	-0.355 – 0.349	1.532
	Two	0.344	-0.008 – 0.622	1.372

ICC = Intra-class correlation coefficient; CI = Confidence interval; SEM = Standard error of measurement

Table 3.4. Test-retest reliability of volatile – stable outcomes for wins and losses on the IBLT.

Outcome	Block	Run	ICC	95% CI	SEM
Wins	Wins volatile – wins stable	One	0.403	0.060 – 0.662	1.367
		Two	0.294	-0.112 – 0.553	1.391
Losses	Losses volatile – losses stable	One	-0.547	-0.753 – -0.240	2.204
		Two	0.473	0.146 – 0.708	0.937

ICC = Intra-class coefficient; CI = Confidence interval; SEM = Standard error of measurement

Table 3.5. Test-retest reliability estimates of the inverse temperature (β) across IBLT task blocks.

Block	Run	ICC	95% CI	SEM
Wins volatile	One	0.355	0.005 – 0.630	0.811
	Two	0.434	0.098 – 0.683	1.090
Losses volatile	One	0.094	-0.266 – 0.433	1.182
	Two	0.441	0.106 – 0.687	1.009
Both volatile	One	0.552	0.248 – 0.757	0.837
	Two	0.551	0.247 – 0.757	0.844

ICC = Intra-class correlation coefficient; CI = Confidence interval; SEM = Standard error of measurement

Table 3.6. Test-retest reliability estimates of the proportion of win-driven choices across IBLT task blocks.

Block	Run	ICC	95% CI	SEM
Wins volatile	One	0.479	0.154 – 0.712	0.080
	Two	0.340	-0.013 – 0.619	0.080
Losses volatile	One	-0.026	-0.374 – 0.330	0.167
	Two	0.564	0.265 – 0.765	0.077
Both volatile	One	-0.366	-0.636 – -0.015	0.200
	Two	0.173	-0.190 – 0.496	0.137

ICC = Intra-class correlation coefficient; CI = Confidence interval; SEM = Standard error of measurement

3.3.3. Practice effects

3.3.3.1. Learning rates

Finally, I assessed potential practice effects on learning parameters in the Information Bias Learning Task. There was a main effect of Session ($F(1,29) = 4.42$, $p < 0.05$, $\eta^2 = 0.006$), as well as a Block x Session interaction for the learning rates ($F(2,58) = 3.80$, $p < 0.05$, $\eta^2 = 0.008$). As shown in Figure 3.6.A, learning rates were lower in Session 1 than in Session 2 for the Loss-volatile blocks ($t(119) = -3.60$, $p < 0.001$), while there was no difference between the sessions for the Win-volatile blocks ($t(119) = -0.91$, $p = 0.364$) or Both-volatile blocks ($t(119) = 0.05$, $p = 0.958$). In addition, there was a Block x Run interaction ($F(2,58) = 4.73$, $p < 0.05$, $\eta^2 = 0.009$), with learning rates being higher in Run 1 than in Run 2 (see Figure 3.6.B) for the Both-volatile blocks ($t(119) = 2.59$, $p < 0.05$). In contrast,

there were no significant differences between Run 1 and Run 2 for the Win-volatile blocks ($t(119) = -0.83$, $p = 0.409$) or the Loss-volatile blocks ($t(119) = -0.97$, $p = 0.332$). Finally, a four-way Block x Session x Outcome valence x Run interaction ($F(2,58) = 6.62$, $p < 0.01$, $\eta^2 = 0.008$) was observed. As shown in Figure 3.7., this interaction was due to win learning rates being higher than loss learning rates in the first run of the Both-volatile block in Session 1 ($t(29) = 3.05$, $p < 0.01$), whereas there was no difference between win and loss learning rates in the first run of Session 2 ($t(29) = -0.46$, $p = 0.650$). Overall, these findings suggest that healthy participants may initially utilise informative positive outcomes to a greater extent than informative negative outcomes, but this difference disappears with repeated exposure to the task.

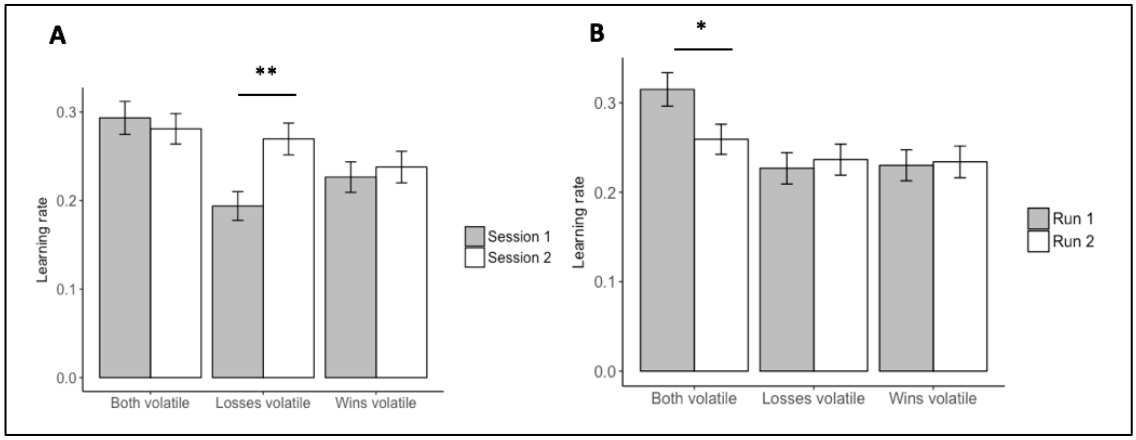

Figure 3.6. (A) Learning rates were higher in Session 2 than in Session 1 during the Loss-volatile block, indicating a potential practice effect across the sessions. **(B)** Learning rates decreased from Run 1 to Run 2 for the Both-volatile blocks within each session. This may reflect an effect of fatigue on task performance over time, as for Both-volatile blocks Run 1 represents the first block and Run 2 the last block that participants completed on the task. Error bars represent the standard error of the mean (*$p < 0.05$, **$p < 0.001$).

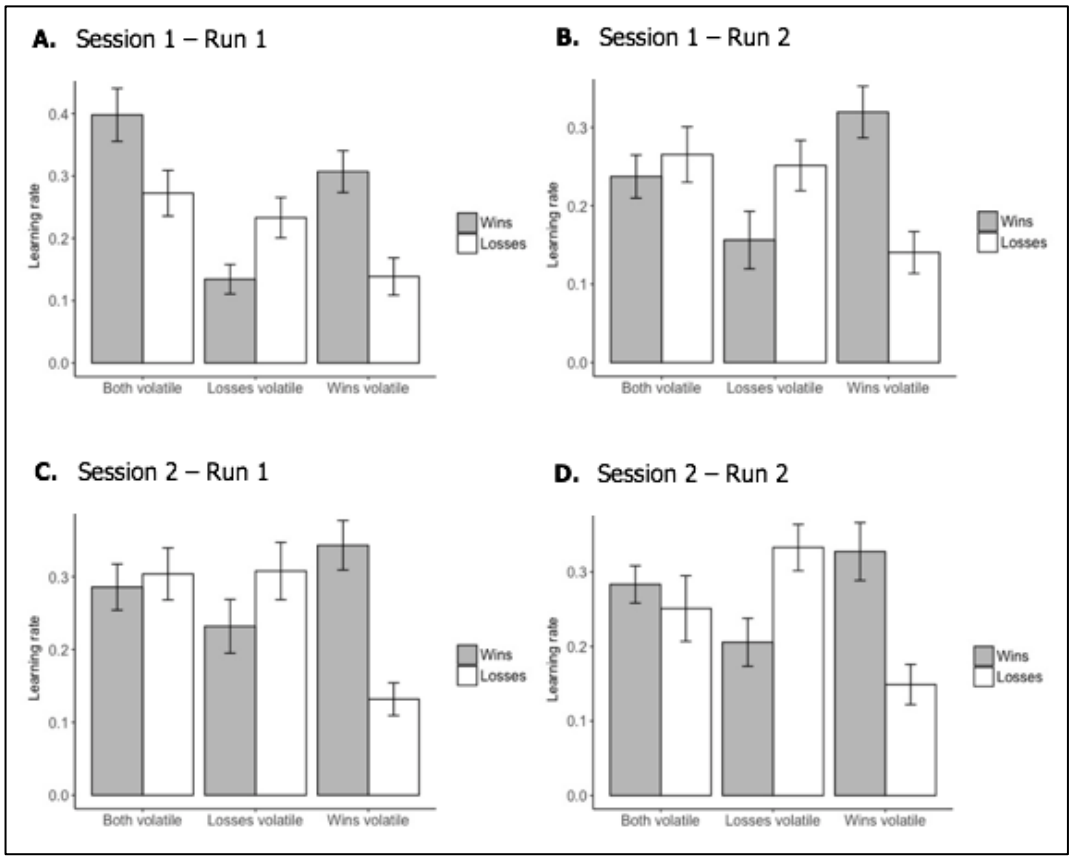

Figure 3.7. IBLT learning rates by outcome and block type. (A) Learning rates in Run 1 of Session 1. At baseline, learning rates were higher for win outcomes than loss outcomes in the Both-volatile block. (B) Learning rates in Run 2 of Session 1. (C) Learning rates in Run 1 of Session 2. (D) Learning rates in Run 2 of Session 2. Error bars represent the standard error of the mean.

3.3.3.2. Inverse temperature

For the inverse temperature, no main effects of Block, Session, or Time were found (all p >0.05). However, as shown in Figure 3.8., there was a Block x Session interaction ($F(2,58) = 2.26$, $p < 0.05$, $\eta^2 = 0.009$). Specifically, β-values were lower in the Loss-volatile block than the Both-volatile block ($t(59) = 2.57$, $p < 0.05$) and Win-volatile block ($t(59) = 2.59$, $p < 0.05$) in Session 2, whereas in Session 1 there were no differences between the Loss-volatile and Both-volatile ($t(59) = -0.13$, $p = 0.899$) or Win-volatile block ($t(59) = -0.83$, $p = 0.411$). It therefore appeared that participants' choices were more random in the Loss-volatile block than the other two blocks in the second testing session. However, there was

no significant difference in choice stochasticity between Session 1 and Session 2 for any of the task blocks (all p >0.05), suggesting that there was no change in randomness of choices within the individual blocks over time.

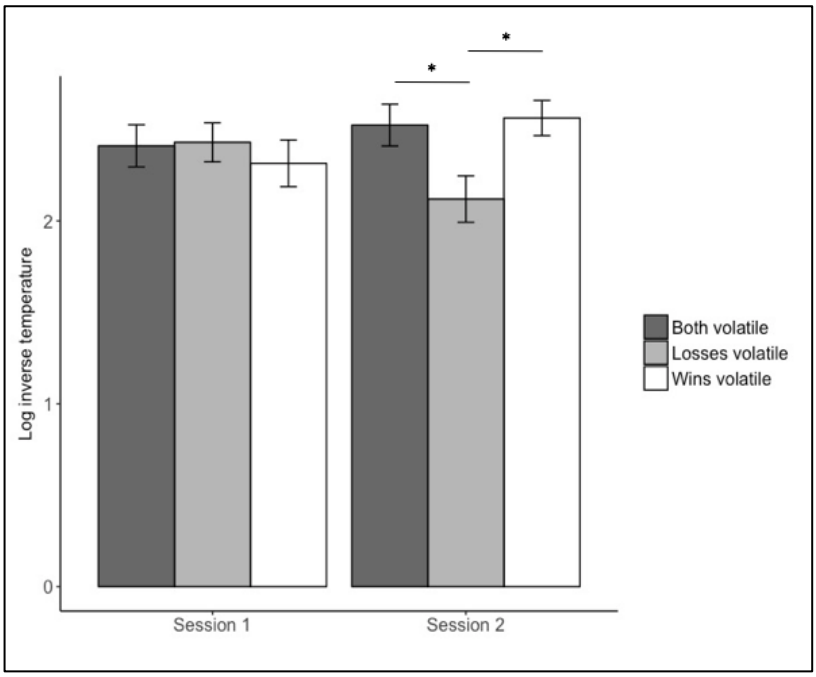

Figure 3.8. Log inverse temperature values by Session and Block on the IBLT. Error bars represent the standard error of the mean (*p <0.01).

3.3.3.3. Win-driven choice behaviour

For win-driven choice behaviour, there was a main effect of Block (F(2,58) = 87.016, p <0.001, η^2 = 0.499), such that the proportion of win-driven choices was greater in the Win-volatile blocks than in the Both-volatile (t(119) = 9.23, p <0.001) and Loss-volatile blocks (t(119) = 17.79, p <0.001). In addition, participants showed greater win-driven choice behaviour in the Both-volatile blocks than the Loss-volatile blocks (t(119) = 7.16, p <0.001). This suggests that participants based their choices mainly on win outcomes when only wins were volatile, with intermediate win-driven choice behaviour when both wins and loss were volatile, and fewest win-driven choices when wins were stable. There was also a main effect of Run (F(1,29) = 6.86, p <0.05, η^2 = 0.011), with participants making more win-driven choices in Run 1 (M = 0.54, SE = 0.01) than in Run 2 (M = 0.51, SE =

0.01).

Furthermore, there was a Session x Run interaction ($F(1,29) = 7.86$, $p < 0.01$, $\eta^2 = 0.009$). A three-way Block x Session x Run interaction ($F(2,58) = 7.96$, $p < 0.001$, $\eta^2 = 0.044$) indicated that this effect differed across the task blocks. When running separate analyses for the individual blocks, it was found that this interaction was significant for the Both-volatile blocks ($F(1,29) = 11.10$, $p < 0.01$, $\eta^2 = 0.09$) but not the Win-volatile ($F(1,29) = 3.69$, $p = 0.646$) or Loss-volatile blocks ($F(1,29) = 3.93$, $p = 0.569$). As depicted in Figure 3.9., win-driven choice behaviour decreased in the Both-volatile block from Session 1 to Session 2 in Run 1 ($t(29) = -2.40$, $p < 0.05$), whereas it increased in Run 2 ($t(29) = 2.34$, $p < 0.05$). In addition, the proportion of win-driven choices for the Both-volatile block was higher in Run 1 than in Run 2 during Session 1 ($t(29) = 3.19$, $p < 0.01$), whereas the opposite pattern was found in Session 2 ($t(29) = -2.34$, $p < 0.05$). It therefore appeared that win-driven behaviour was relatively changeable across runs and sessions in blocks where both win and loss outcomes were volatile.

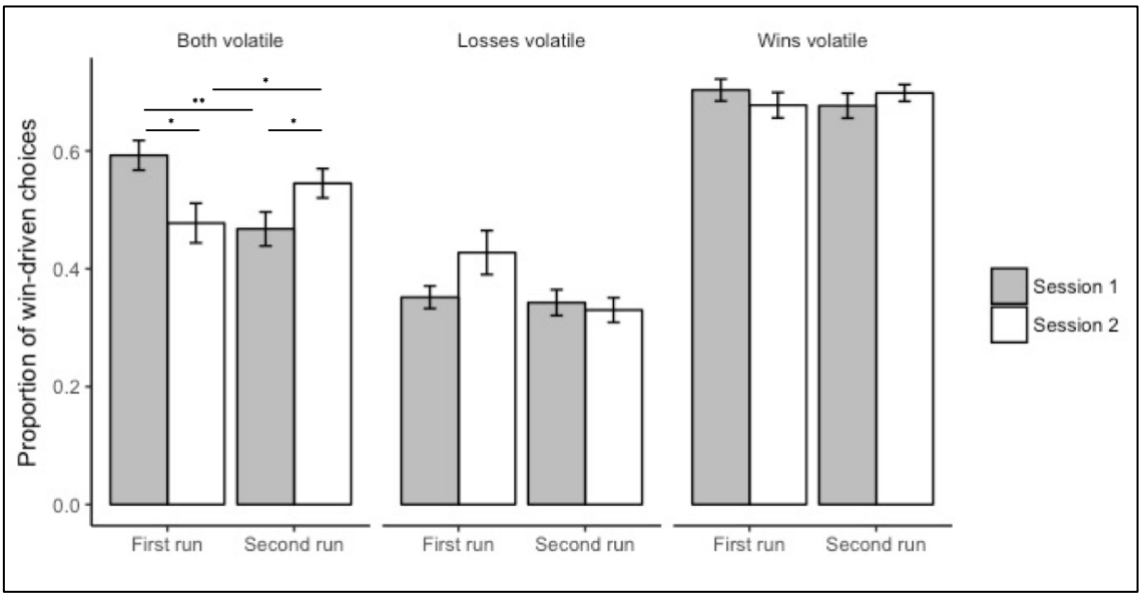

Figure 3.9. Proportion of win-driven choices on the IBLT by Block, Session, and Run. Error bars represent the standard error of the mean (* $p < 0.05$; ** $p < 0.01$).

3.4. Discussion

3.4.1. IBLT performance

In this study, I examined overall performance, test-retest reliability, and practice effects on the Information Bias Learning Task in healthy adults. First of all, reports by Pulcu and Browning (2017) that healthy individuals can generate separate information content estimates for positive and negative outcomes as reflected in learning rates were replicated. Furthermore, participants were generally able to adjust their learning rates in response to volatility manipulations of the environment, with learning rates being higher for changeable than stable associations.

3.4.2. Test-retest reliability

In addition, test-retest reliability of the IBLT was investigated to explore consistency of task performance across measurements. This is the first study to examine the stability of computational parameter values in a cognitive task over multiple testing sessions. Intraclass correlation coefficients for learning rates, inverse temperature values, and win-driven choice behaviour ranged from poor to moderate depending on the task block and outcome valence. For instance, poor ICC values were reported for win learning rates in the Both-volatile blocks, while win learning rates in the Win-volatile blocks showed moderate retest reliability. Overall, ICC values for learning rates tended to be higher for volatile outcomes than for stable outcomes. This may indicate that participants had more variable responses to uninformative events than to informative events across the sessions.

There are several possible explanations for the inconsistency in task performance observed here. Firstly, it is possible that the underlying construct being measured with the IBLT is changeable. That is, learning rates for positive and negative events might reflect a variable state rather than a stable information processing bias of the participant. While the IBLT may not be suitable as a diagnostic tool if learning rates indicate a state rather

than a trait variable, the task could be useful for investigating whether specific interventions can elicit a positive state (i.e., increased learning from positive outcomes). Secondly, it has been suggested that low ICC values can be caused by minimal between-subjects variability rather than large within-subjects variability (Weir, 2005). That is, if learning rates tend to be homogenous across participants, minimal changes in task performance over time will lead to low ICC values. This appears to be a less likely explanation for the present findings of low reliability, however, as learning rates ranged from 0.01 to 0.97 for volatile outcomes and 0.01 to 0.89 for stable outcomes across participants. Thirdly, it is possible that performance is especially variable in the first session, where participants are first presented with the task, and becomes more stable with more extensive experience on the task. It would therefore be of interest to investigate whether retest reliability is higher across a greater number of testing sessions. Finally, it should be emphasised that, while ICCs are useful for providing an indication of variability in task performance over time, these values should not be overinterpreted in isolation. It has been suggested that the interpretation of ICC scores can be facilitated with the additional use of SEM scores, which provide an indication of the variation in measurement errors when the true score is held constant (Dudke, 1979; Harvill, 1991). In other words, if the observed change in task performance is greater than the SEM value, this suggests that the difference reflects a true change in the measured construct. The SEM values reported here may therefore be of use in interpreting alterations in IBLT performance in repeated-measurement studies.

Overall, test-retest reliability thus appears to be suboptimal on the Information Bias Learning Task, which could result in an increased risk of Type II errors in repeated-measures studies. While this may indicate that the IBLT is less appropriate for measuring learning changes in longitudinal studies, the findings do not necessarily indicate that this

task cannot be of use in other research designs. First of all, the present data were collected in the absence of any interventions. The results therefore do not provide information on whether effects of specific interventions on IBLT performance are consistent over time. In addition, it is possible that the IBLT can be used to modify rather than assess learning processes. For example, it has been shown that a commonly used measure of attentional biases called the dot-probe task, which has poor test-retest reliability (Cooper et al., 2011; Schmukle, 2005; Staugaard, 2009), can improve clinical symptoms of anxiety disorders when applied as a training paradigm (Amir et al., 2009; Linetzky et al., 2015; Schmidt et al., 2009). Similarly, the IBLT may have potential as a therapeutic tool in psychiatric disorders by modulating affective learning biases. Additional research is needed to explore these applications of the IBLT in both healthy adults and psychiatric populations.

3.4.3. Practice effects

Finally, I examined potential practice effects on the IBLT across the two testing sessions. It was found that learning rates increased for the Loss-volatile blocks across the sessions. Although the increase in loss learning rates may be due to enhanced understanding of the task over time, there was also an increase in win learning for this block despite the fact that wins were uninformative. For the Win-volatile and Both-volatile blocks, no such changes in learning rates were observed over time. Moreover, learning rates in these blocks were higher overall than in the Loss-volatile blocks. These findings suggest that participants may have had more difficulties in performing Loss-volatile than Win-volatile or Both-volatile blocks, potentially due to an inability to use negative information to the same extent as positive information in adjusting their performance. Intervention studies should therefore be aware of the potential for increases in loss learning rates across sessions independent of the treatment conditions. Such effects may be accounted for by randomising the order of interventions and control conditions and/or introducing a practice

session prior to the treatment of interest. Further support for an initial bias towards learning from win outcomes rather than loss outcomes was provided by the finding that win learning rates were higher than loss learning rates in the Both-volatile block in Session 1, while the two outcomes are in reality equally informative. In the second session there were no significant differences between learning rates for the two outcomes in this block, indicating that participants had acquired a balance in the use of win and loss outcomes. Finally, a decrease in learning rates for the Both-volatile blocks was observed from Run 1 to Run 2. This may reflect an effect of fatigue over time, as the first Both-volatile block was presented at the start of the task, whereas the second Both-volatile block was performed at the very end of the task. For the inverse temperature, values were shown to be similar across the blocks in Session 1, whereas there was an indication that choice stochasticity was greater in the Loss-volatile blocks than the Win- or Both-volatile blocks in Session 2. However, there was no significant difference in randomness of choice across the two sessions for any of the blocks. Overall, randomness of choice therefore did not appear to change for the individual IBLT blocks with repeated measurements.

Finally, win-driven choice behaviour was found to be relatively consistent in the Win- and Loss-volatile blocks, whereas performance of the Both-volatile blocks was variable across measurements. Win-driven choices increased across the task runs in the first session, whereas the opposite pattern was observed in the second session. The high win-driven behaviour in the first run of the first session could reflect a similar initial bias towards learning from positive outcomes as observed for the learning rates. The reduction in win-driven choices in the second run might therefore indicate a better understanding of the information content of the win and loss outcomes. It is unclear, however, why this pattern was reversed in the second testing session. One possible explanation is that the equal information content of the win and loss outcomes resulted in alternating between

decisions based on positive or negative events. Alternatively, tiredness may have been a source of variability in performance of these blocks across the two runs, which were carried out at the start and the very end of the testing sessions.

3.4.4. Conclusion

This study has replicated previous findings that healthy adults can adapt their learning rates for positive and negative outcomes individually in response to changes in the volatility of the environment. Moreover, it has provided an indication of the expected contribution of errors to changes in scores on the IBLT over time. Finally, it was shown that participants may have initial difficulties in attending to negative outcomes when adjusting their behaviour on the task, but can acquire this ability with practice. In the future, these practice effects should therefore be taken into account when interpreting changes in performance on the IBLT with repeated measurements.

Chapter 4. Increasing learning rates for positive outcomes with bifrontal tDCS: A proof-of-concept study

4.1. Introduction

Cognitive neuropsychological models suggest that information processing biases play a crucial role in the development and maintenance of depressive disorders. Beck's psychological theory of depression proposes that adverse events in early life can lead to the construction of negative schemas, or pervasive patterns of thought and behaviour. In later life, such schemas can be activated and reinforced through the experience of additional stressors. In turn, this can lead to systematic biases towards processing of and attending to negative information relative to positive information (Beck, 1967; Beck, 2008). In support of this theory, evidence from experimental studies has shown that individuals prone to mood disorders preferentially attend to aversive cues, recall more unpleasant memories, and interpret ambiguous events in a more negative manner than healthy controls (Mathews & Macleod, 2005; Peckham, McHugh, & Otto, 2010). Importantly, it has been suggested that attending to negative stimuli can prevent the processing of other, potentially more positively valenced events (Disner et al., 2011). A negative bias could thus conceivably induce a vicious circle, where attending to negative information increases low mood, which can consequently strengthen core biases and negative beliefs. Attentional biases towards positive stimuli, in contrast, have been found to be correlated with higher trait resilience to stress (Thoern et al., 2016). There has therefore been increasing interest in investigating the possibility of altering attentional biases in people with or at risk of depression.

For example, transcranial direct current stimulation (tDCS) studies have suggested that anodal stimulation of the left dorsolateral prefrontal cortex (DLPFC) may reduce attention to negative information (Clarke et al., 2014; Heeren et al., 2015; Heeren et al.,

2017; Ironside et al., 2015), although this research has tended to focus on biases associated with anxiety rather than depression. It is presently unclear whether depression-related affective biases can similarly be altered with tDCS, although research in healthy adults has suggested that this intervention may produce more pronounced effects on biased information processing associated with anxiety than depression (Ironside et al., 2015). Of interest, similar indications of differential effects on anxiety and depressive symptoms have been observed with alternative interventions targeting attentional biases. For example, various studies have attempted to modify biases in attention and information processing with the use of cognitive training paradigms (Browning, Holmes, & Harmer, 2010; Browning et al., 2012; Hakamata et al., 2010). However, while such interventions have again been applied with some success in anxiety disorders, meta-analyses have indicated that their impact on depressive symptoms tends to be minimal and short-lasting (Hallion & Ruscio, 2011; Jones & Sharpe, 2017). A potential explanation for this small effect size may be that these training paradigms can lead to immediate avoidance of negative stimuli (Hallion & Ruscio, 2011), but fail to change underlying beliefs. An interesting alternative approach may therefore be to explore whether it is possible to modify the *cause* of attentional biases associated with depression.

Recently, it has been proposed that negative affective biases may be due to a belief that negative events are more *informative* than positive events (Pulcu & Browning, 2017). Theoretically, it should therefore be possible to reduce such biases by altering people's information content estimates of positive and negative outcomes. As outlined in Section 2.2., it has been shown that these estimates as reflected by learning rates can be manipulated in a reinforcement learning task named the Information Bias Learning Task (IBLT). By varying the volatility and thereby the information content of win and loss outcomes, this task influences how quickly individuals learn from positive or negative

events (Browning et al., 2015; Pulcu & Browning, 2017). It would be of interest to investigate whether learning rates on this task can be further modulated with the application of non-invasive brain stimulation. Previously, Hämmerer et al. (2016) showed that anodal stimulation of the left dorsolateral prefrontal cortex (DLPFC) with the cathode below the inion over the occipital lobe in healthy adults slightly increased global learning rates in a value-based decision making task. In patients with schizophrenia, anodal stimulation over the medial prefrontal cortex was similarly found to enhance and normalise learning rates (Reinhart et al., 2015). There is thus preliminary evidence suggesting that modulation of prefrontal activity can affect learning rates. However, the paradigms used in these two studies did not allow for measuring differences in responses to rewarding and aversive outcomes, leaving open the question whether tDCS of the DLPFC can specifically modify learning from positive or negative information.

In this proof-of-concept study, I tested whether the tDCS protocol described in Section 2.4. could change learning rates in any direction for win and loss outcomes on the IBLT in healthy adults. It was expected that bifrontal stimulation of the DLPFC would not only affect activity in the prefrontal cortex, but also in the connected dorsal anterior cingulate cortex (dACC), which has been shown to be linked to volatility estimates and associated learning rates (Behrens et al., 2007). Since the dACC response to volatility changes has thus far only been described for rewarding outcomes (Behrens et al., 2007; Behrens et al., 2008), it was unclear whether bifrontal tDCS might modulate learning rates for negative outcomes as well. Based on initial findings from an EEG study showing that prefrontal tDCS specifically influences ACC activity for positive but not negative outcomes (Vanderhasselt et al., 2013), however, it was hypothesised that win learning rates may be altered more than loss learning rates. In addition, the use of volatile and stable associations provided an opportunity to test the effects of tDCS in circumstances where

optimal decisions should involve high and low learning rates, respectively. As volatile associations tend to lead to high learning rates, it could be expected that increases in the speed of learning with stimulation in this condition might not be readily observed due to potential ceiling effects. In stable conditions, on the other hand, there might be greater scope for detecting higher learning rates with tDCS. I also investigated possible changes in randomness of choices as measured with inverse temperature values, as there are some indications that this parameter might also be changed with prefrontal stimulation (see Chapter 1; Turi et al., 2015).

Furthermore, I compared the effects of tDCS applied during task performance ('online' tDCS) with tDCS applied prior to task performance at rest ('offline' tDCS). It has been suggested that non-invasive brain stimulation is particularly effective at influencing neurocognitive processing if the stimulation is applied while the targeted brain region is actively engaged (Dedoncker et al., 2016; Martin et al., 2014). Furthermore, tDCS applied during task performance may strengthen long-term potentiation (LTP) effects of stimulation as described in Chapter 2, and thereby enhance learning processes (Woods et al., 2016). It was therefore expected that online tDCS would increase learning rates to a greater extent during and after stimulation compared with offline tDCS.

4.2. Methods and Materials

4.2.1. Participants

Ethical approval was obtained from the University of Oxford Central University Research Ethics Committee (CUREC; R48995/RE001). For the online tDCS study, twenty healthy participants aged 18 to 45 years (10 female, mean age = 25.00, SD ± 4.22 years) were recruited from the community via advertisements. A new sample of twenty healthy adults was recruited for the offline tDCS study (15 female, mean age = 25.15, SD ± 5.81 years). In addition, based on the findings from these two studies, I conducted a replication

experiment of the online tDCS study with a different sample of 20 healthy volunteers (7 female, mean age = 24.15, SD ± 3.62). Written informed consent was given by all participants prior to taking part in the testing sessions. Participants were reimbursed for their time and travel expenses at a rate of £10 per hour and received all the money they won in the computerised task. Exclusion criteria were a history of psychiatric disorders, neurological illness, use of psychoactive medication, personal or family history of epileptic fits or seizures, and any contraindications to tDCS.

4.2.2. General procedure

All participants were invited to two stimulation sessions, which took place at least one week apart to prevent carry-over effects. In one of the sessions participants received active bifrontal tDCS, whereas in the other session they received sham tDCS. Order of the tDCS conditions was randomised and double-blinded using a study mode on the tDCS device. Following provision of written informed consent, the researcher conducted a screening procedure to ensure safety of tDCS application (see Appendix D). If the screening was successful, participants were asked to complete Beck's Depression Inventory II (BDI; Beck, Steer, & Brown, 1996) and the Trait subscale of the State-Trait Anxiety Inventory (STAI-Trait; Spielberger et al., 1983). These questionnaires were included to assess potential correlations between symptoms of depression or anxiety and reinforcement learning parameters (see Appendix B). To investigate state mood and anxiety, participants were given the Positive and Negative Affect Scales (PANAS; Watson, Clark, & Tellegen, 1988) and the State subscale of the State-Trait Anxiety Inventory (STAI-State; Spielberger et al., 1983). Next, the researcher explained the Information Bias Learning Task (IBLT; described in Chapter 3), which was presented on a laptop monitor using *Presentation* software 18.3 (Neurobehavioral Systems Inc, Albany, CA). The tDCS equipment was then set up and briefly tested to ensure that the participant was comfortable with the stimulation.

All participants were asked to complete one Both-volatile block of the IBLT without stimulation to gain a baseline measurement of task performance when both win and loss outcomes were informative (see Figure 4.1.). In the online tDCS study, participants then received 20 minutes of stimulation while completing a Win-volatile and a Loss-volatile task block. After the stimulation had ended, participants were again asked to perform a Win-volatile and a Loss-volatile block, followed by a final Both-volatile block. Order of the Win-volatile and Loss-volatile blocks was counterbalanced across participants, with individual participants receiving the same block order across the two tDCS sessions. In the offline tDCS study, participants completed the first Both-volatile block, which was followed by 20 minutes of tDCS received at rest. After the stimulation had ended, participants were asked to complete the remaining five blocks of the IBLT. Once the task had been completed, participants again filled out the STAI-State and PANAS questionnaires to assess potential changes in acute mood and anxiety. These questionnaires were included as a control measure, as previous studies have found no evidence that a single session of bifrontal tDCS affects mood in healthy adults (Morgan et al., 2014; Plazier et al., 2012; Vanderhasselt et al., 2013). At the end of each session, a questionnaire was used to assess potential adverse effects of tDCS experienced by the participant (see Appendix E). In addition, participants were asked to indicate whether they believed they had received active or sham tDCS in each session to determine whether participant blinding was effective.

Condition	**5 min**	**15 min**	**~5 min**	**20 min**	**~15 min**	**5 min**
Online tDCS	Questionnaires	Task description and tDCS preparation	IBLT block 1	tDCS IBLT block 2-3	IBLT block 4-6	Questionnaires
Condition	**5 min**	**15 min**	**~5 min**	**20 min**	**~30 min**	**5 min**
Offline tDCS	Questionnaires	Task description and tDCS preparation	IBLT block 1	tDCS at rest	IBLT block 2-6	Questionnaires

Figure 4.1. Timelines of the general procedure, tDCS application, and behavioural tasks in the online and offline tDCS studies.

4.2.3. Transcranial direct current stimulation (tDCS)

Stimulation was delivered using a battery-powered device (Eldith DC-Stimulator Plus, Neuroconn, Germany). The rubber electrodes (5 x 5 cm) were placed in saline-soaked sponges and attached to the scalp using rubber bands. In both sessions, the anodal electrode was placed on the left DLPFC, while the cathodal electrode was placed over the right DLPFC (F3 and F4, respectively, according to the 10/20 system of electrode placement). In the active tDCS condition, stimulation was delivered at 2 mA for 20 minutes, with 10s ramping-up and ramping-down. In the sham tDCS condition, participants received 30s of direct current followed by impedance control with a small current pulse being produced every 550 ms (110 μA over 15 ms), resulting in an instantaneous current of no more than 2 μA. The sham stimulation procedure ensured that participants experienced the physical sensations associated with active tDCS. Furthermore, this method allowed for effective double blinding by showing realistic impedance values on the device display. In the online tDCS condition, stimulation was applied during completion of IBLT block 2 and 3, with Blocks 4 – 6 serving to assess whether potential effects of tDCS persisted after the

stimulation had ended. In the offline tDCS condition, blocks 2 – 6 were used to examine changes in reinforcement learning following stimulation at rest.

4.2.4. Statistical analyses

4.2.4.1. Learning parameters

The effects of tDCS on learning parameters across task blocks were examined with repeated-measures ANOVAs using R software (R Version 3.3.1.). For the Win- and Loss-volatile blocks carried out during and after stimulation, Block (Win-volatile versus Loss-volatile), Outcome valence (win outcome versus loss outcome), Run (first 2 blocks versus second 2 blocks), and tDCS condition (sham tDCS versus bifrontal tDCS) were included as within-subject factors, with learning rate as the main outcome variable. For the Both-volatile blocks, Time (before tDCS versus after tDCS), Outcome valence, and tDCS condition were used as within-subject factors. Changes in the inverse temperature were investigated in repeated-measures ANOVAs similar to the learning rates, with exclusion of the outcome valence variable as a single inverse temperature value was calculated for both win and loss outcomes. Order of tDCS administration (sham tDCS first versus active tDCS first) was included as a between-subjects factor in all analyses to account for potential differences due to stimulation order.

Significant interactions were followed up with separate repeated-measures ANOVAs and/or paired *t*-tests. An inverse logit transform was applied to the learning rates in order to limit the influence of extreme outlier values. For the inverse temperature values a log transform was used to enhance normality of the data. To investigate potential influences of baseline learning rates on tDCS effects, statistical analyses were run both with and without baseline win – loss learning rate as a covariate. Figures and reported values represent raw parameter values to facilitate interpretation of the results.

As there have been some indications that non-invasive brain stimulation may affect

model-based learning (Smittenaar et al., 2014; Wittkuhn et al., 2018), I also examined whether there was a change in the fit of the model-based learning (Model 1) versus model-free learning (Model 3) models as described in Chapter 3 with active compared to sham tDCS. As no significant effect of tDCS was found for either the present or follow-up studies (Chapters 4-6), however, results from these analyses are reported in Appendix C.

4.2.4.2. Win-driven choice behaviour

In addition to computationally derived parameters, changes in win-driven choice behaviour as described in Section 3.2.5. were assessed in repeated-measures ANOVAs, with the proportion of win-driven choices as the outcome variable and tDCS condition, Block, Run/Time, and tDCS order as independent variables. Simple behavioural changes were also investigated by comparing the amount of money won in sham versus active tDCS sessions using paired *t*-tests.

4.2.4.3. Mood and anxiety measures

Finally, I examined associations between bifrontal tDCS and changes in acute mood and anxiety using PANAS and STAI-State questionnaire scores. I applied a repeated-measures ANOVA with PANAS Positive score, PANAS Negative score, and STAI-State score as the outcome variables, and Time (pre- versus post-stimulation), tDCS condition, and tDCS order as the predictor variables.

4.3. Results

4.3.1. Participant characteristics

The online and offline tDCS groups did not differ in terms of gender, age, or baseline (pre-stimulation) mood and anxiety scores (see Table 4.1.). In the sham tDCS condition, there was a significant difference between the online and offline stimulation group in baseline PANAS Negative Affect scores, which were higher in the offline group (M = 12.95, SE = 0.66) than the online group (M = 10.90, SE = 0.25). Within the groups, however, there

were no differences in baseline PANAS or STAI-State scores between the sham tDCS and bifrontal tDCS sessions (all $p > 0.05$). This suggests that any observed differences in the effects of sham versus active stimulation were not likely to reflect differences in pre-stimulation mood and anxiety state between the two sessions.

To assess potential baseline biases towards learning from either positive or negative information prior to tDCS, I compared learning rates for win and loss outcomes in the first block (Both-volatile) of the first session of the study. Before receiving either sham or bifrontal tDCS in the first session, participants in the online tDCS group had a higher learning rate for win outcomes than loss outcomes in the Both-volatile block ($t(19) = 2.33$, $p < 0.05$). Notably, win and loss learning rates were similar in the first (pre-stimulation) block across the two sessions (Outcome valence x tDCS condition: $F(1,18) = 1.02$, $p = 0.326$). In the offline tDCS group there was no significant difference in learning rates for win and loss outcomes in the Both-volatile block of the first session ($t(19) = -0.09$, $p = 0.929$). Similar to the online tDCS group, there was no difference in the baseline win and loss learning rates between the sham and active tDCS sessions ($F(1,18) = 0.21$, $p = 0.656$) in the Both-volatile block before stimulation. Hence, it was improbable that any effects of bifrontal versus sham tDCS on learning rates were caused by pre-stimulation differences in win and loss learning rates.

Table 4.1. Mean (SE) baseline trait and state anxiety and mood questionnaire scores by tDCS group. *t*-values and *p*-values represent the difference between scores in the online and offline tDCS groups.

Questionnaire	Online tDCS	Offline tDCS	*t*-value	*p*-value
BDI	5.11 (1.20)	5.55 (1.79)	-0.21	0.84
STAI-Trait	37.47 (2.05)	38.45 (1.77)	-0.36	0.72
PANAS Positive				
Sham tDCS session	26.85 (1.09)	28.15 (1.43)	-0.72	0.48
Bifrontal tDCS session	27.20 (1.19)	28.65 (1.41)	-0.79	0.44
PANAS Negative				
Sham tDCS session	10.90 (0.25)	12.95 (0.66)	-2.91	0.01
Bifrontal tDCS session	11.65 (0.50)	12.00 (0.60)	-0.45	0.66
STAI-State				
Sham tDCS session	29.95 (1.04)	32.25 (1.94)	-1.04	0.31
Bifrontal tDCS session	30.00 (1.38)	29.55 (1.64)	0.21	0.84

BDI; Beck's Depression Inventory II, PANAS; Positive and Negative Affect Scale; STAI, State-Trait Anxiety Inventory.

4.3.2. Study 1. Effects of online tDCS on reward learning in the IBLT

4.3.2.1. Learning rate Win-volatile and Loss-volatile blocks

The main outcome of interest in the present study was a potential change in learning parameters due to bifrontal tDCS. First, I explored whether tDCS condition affected learning rates on the Win-volatile and Loss-volatile task blocks across the sessions. As expected, there was a significant two-way Block x Outcome valence interaction ($F(1,18) = 94.92$, $p < 0.001$, $\eta^2 = 0.162$). Win learning rates were higher in the Win-volatile blocks compared to the Loss-volatile blocks ($t(79) = 7.02$, $p < 0.001$), whereas loss learning rates were higher in the Loss-volatile than in the Win-volatile blocks ($t(79) = 8.35$, $p < 0.001$). In addition, learning rates for win outcomes were higher than for loss outcomes in the Win-volatile blocks ($t(79) = 10.02$, $p < 0.001$), whereas the opposite pattern was observed in Loss-volatile blocks ($t(79) = 5.23$, $p < 0.001$). Overall, this indicated that participants tended to have higher learning rates for volatile outcomes than for stable outcomes.

A two-way tDCS condition x Block interaction ($F(1,18) = 8.09$, $p < 0.01$, $\eta^2 =$

0.013) was observed. It was found that learning rates were higher in the Loss-volatile blocks with active tDCS than sham tDCS ($t(79) = 3.70$, $p < 0.001$), whereas there was no difference between the two conditions in the Win-volatile blocks ($t(79) = -0.17$, $p = 0.866$). Crucially, there was also a tDCS condition x Outcome valence interaction ($F(1,18) = 4.89$, $p < 0.05$, $\eta^2 = 0.006$). As shown in Figure 4.2., bifrontal tDCS was associated with an increase in win learning rates ($t(39) = 2.77$, $p < 0.01$) but not loss learning rates ($t(39) = 0.50$, $p = 0.618$) compared to sham tDCS. In addition, learning rates for win outcomes were significantly higher than learning rates for loss outcomes with active tDCS ($t(39) = 2.17$, $p < 0.05$) but not sham tDCS ($t(39) = 0.41$, $p = 0.685$). This indicates that bifrontal tDCS had a valence-specific effect on learning rates, with learning rates increasing for win but not loss outcomes. Results were unchanged when win-loss baseline learning rates were included in the analyses as a covariate.

Finally, a main effect of Run was observed ($F(1,18) = 5.55$, $p < 0.05$, $\eta^2 = 0.003$), with learning rates being higher in the first run (M = 0.32, SE = 0.02) than in the second run (M = 0.30, SE = 0.02). This may reflect an effect of fatigue over time, as the full task took approximately 35 – 40 minutes to complete. To assess differences in learning rates across time further, separate analyses were carried out for the two individual runs. During tDCS (Run 1), the tDCS condition x Block interaction was significant ($F(1,18) = 9.41$, $p < 0.01$, $\eta^2 = 0.022$), whereas the tDCS condition x Outcome valence interaction was not ($F(1,18) = 1.09$, $= 0.310$). Inspection of the tDCS Condition x Block interaction reveals that this effect is likely driven by an increase in win learning rates ($t(19) = 2.03$, $p = 0.057$) but not loss learning rates ($t(19) = 0.62$, $p = 0.541$) in the Loss-volatile blocks. After tDCS (Run 2), the tDCS condition x Outcome valence interaction was significant ($F(1,18) = 6.33$, $p < 0.05$, $\eta^2 = 0.009$). Once stimulation had ended, there was thus an overall increase

in learning rates for win outcomes with bifrontal tDCS, whether informative or uninformative.

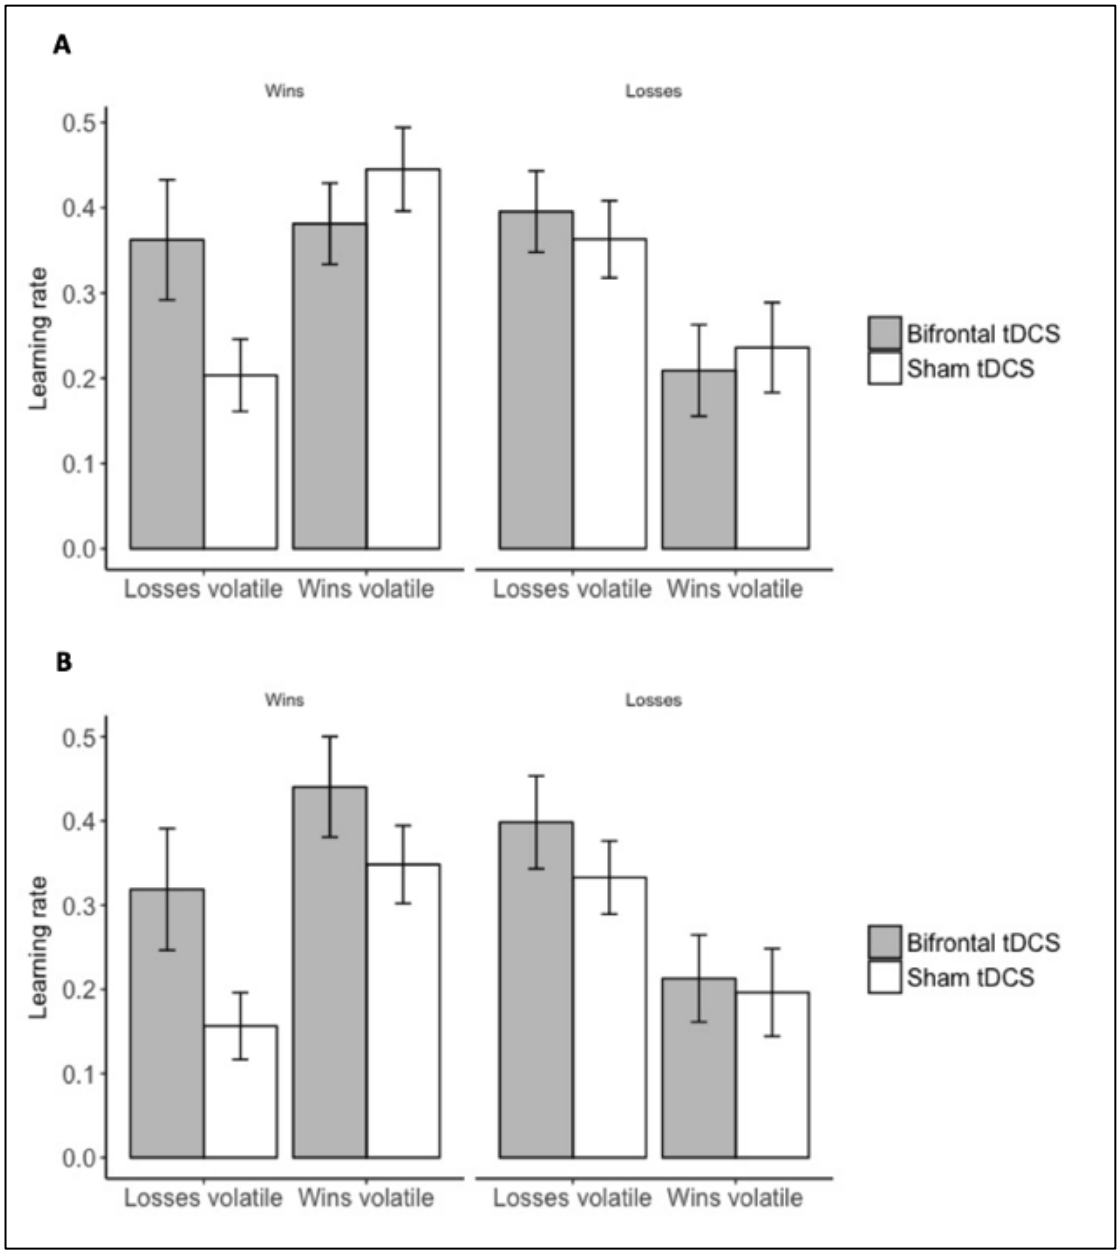

Figure 4.2. Learning rates for win and loss outcomes by tDCS condition for the Win- and Loss-volatile blocks in the online tDCS study. **(A)** Learning rates during tDCS (Run 1). Bifrontal tDCS was associated with higher learning rates in the Loss-volatile block, which appeared to be driven by an increase in win learning rates. **(B)** Learning rates after tDCS (Run 2). Global win but not loss learning rates were higher with bifrontal compared to sham stimulation. Error bars represent the standard error of the mean.

4.3.2.2. Learning rate Both-volatile blocks

Next, I inspected potential differences between pre- and post-stimulation learning parameters with active versus sham tDCS in the Both-volatile blocks. These blocks were mainly included to compare learning rates before and after tDCS and assess whether changes observed in Win- and Loss-volatile blocks during the stimulation transferred to Both-volatile blocks. As an increase in win learning rates was observed in the Win- and Loss-volatile blocks with active tDCS, it was of particular interest to investigate whether there was an increase in win learning rates in the Both-volatile block after compared to before bifrontal stimulation (i.e., a tDCS condition x Outcome valence x Time interaction).

A main effect of Time was found ($F(1,18) = 9.51$, $p < 0.01$, $\eta^2 = 0.036$), with learning rates being higher before (M = 0.41, SE = 0.03) than after active or sham tDCS (M = 0.32, SE = 0.03). In addition, there was a main effect of Outcome valence ($F(1,18) = 5.26$, $p < 0.05$, $\eta^2 = 0.059$), with learning rates being higher for win outcomes (M = 0.42, SE = 0.03) than for loss outcomes (M = 0.31, SE = 0.02). Furthermore, there was a tDCS condition x Outcome valence interaction ($F(1,18) = 5.68$, $p < 0.05$, $\eta^2 = 0.014$). Follow-up analyses indicated that win learning rates were higher than loss learning rates in the bifrontal tDCS condition ($t(39) = 3.85$, $p < 0.001$) but not in the sham tDCS condition ($t(39) = 1.26$, $p = 0.214$). However, although there appeared to be a greater difference between win and loss learning rates after active tDCS compared to sham tDCS (see Figure 4.3.), the tDCS condition x Outcome valence x Time was not significant ($F(1,18) = 0.36$, $p = 0.557$). There is thus no evidence that the increase in win learning rates observed with active tDCS in the Win- and Loss-volatile blocks resulted in changes of win learning rates over time in the Both-volatile blocks.

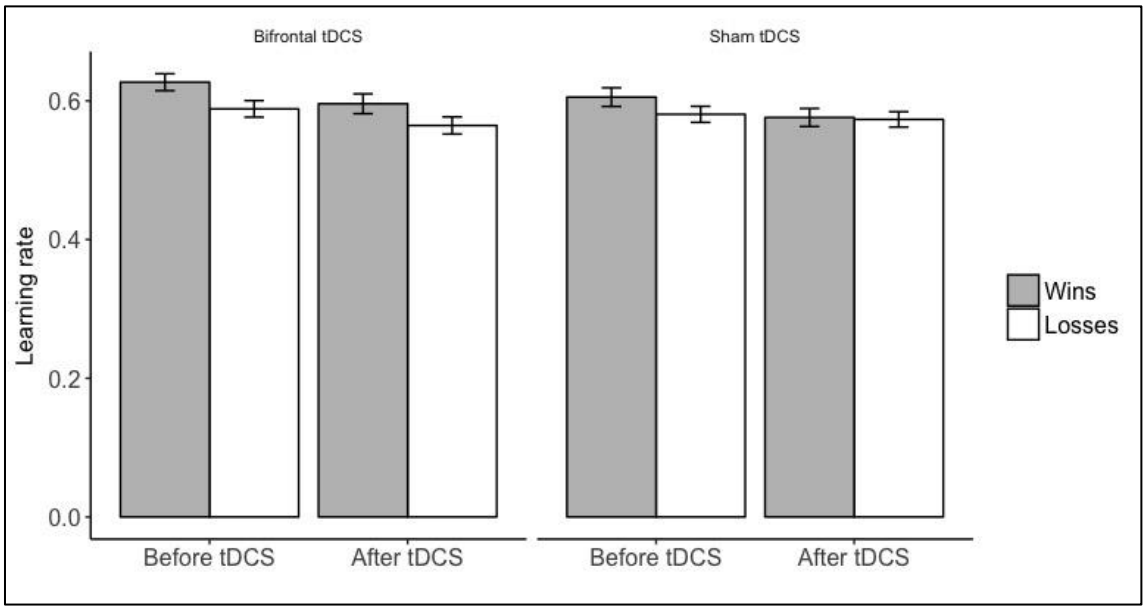

Figure 4.3. Learning rates in the Both-volatile blocks by Outcome valence and tDCS condition over time. Error bars represent the standard error of the mean.

4.3.2.3. Choice stochasticity

In addition to examining changes in learning rates, I investigated whether tDCS affected choice stochasticity as indicated by the inverse temperature values. For the Win-volatile and Loss-volatile blocks, there was no significant main effect of tDCS on the inverse temperature ($F(1,19) = 2.39$, $p = 0.140$). Furthermore, there were no interactions of tDCS condition, Block, and Time in the Win-volatile, Loss-volatile, or Both-volatile blocks (all $p > 0.05$). This suggests that the effects of tDCS were specific to learning rate and did not influence the degree to which participants used expected outcome values to make their choice for a particular shape.

4.3.2.4. Win-driven choice behaviour and total winnings

Next, I investigated whether the proportion of win-driven choices was modified by tDCS condition. Consistent with the findings previously reported by Pulcu and Browning (2017), there was a significant main effect of Block ($F(1,18) = 146.02$, $p < 0.001$, $\eta^2 = 0.748$), with win-driven choices being made more frequently in the Win-volatile blocks (M = 0.69, SE

= 0.01) compared to the Loss-volatile blocks (M = 0.34, SE = 0.01). In addition, there was a significant main effect of tDCS condition ($F(1,18) = 5.30$, $p < 0.05$, $\eta^2 = 0.033$). Importantly, there was an interaction of tDCS condition x Block ($F(1,18) = 7.43$, $p < 0.05$, $\eta^2 = 0.028$), indicating that bifrontal tDCS was associated with a significantly greater number of win-driven choices compared with sham tDCS in the Loss-volatile blocks ($t(39) = 2.18$, $p < 0.05$) but not the Win-volatile blocks ($t(39) = -0.76$, $p = 0.449$). Consistent with the analyses based on formal models, this shows that the effect of tDCS on task performance was more pronounced in the Loss-volatile blocks. For the Both-volatile blocks, there were no significant main or interaction effects (all $p > 0.05$).

Task performance was further investigated by examining differences in money won by the participants across the two tDCS sessions. Paired *t*-tests indicated that these behavioural changes coincided with smaller amounts of total winnings in the active tDCS session than in the sham tDCS session ($t(19) = -2.55$, $p < 0.05$). This suggests that the increase in win-driven choices and learning rates in the Loss-volatile blocks impaired performance on the task in terms of total winnings. That is, by attempting to base decisions on uninformative win outcomes, participants made fewer correct choices and earned less money during performance of the IBLT.

4.3.2.5. Mood and anxiety measures

In addition to examining learning behaviour on the IBLT, I assessed whether acute mood and anxiety measures were altered by the application of bifrontal tDCS. There was a main effect of Time on PANAS Positive scores ($F(1,19) = 21.43$, $p < 0.001$, $\eta^2 = 0.090$), with scores being higher before (M = 27.03, SE = 0.80) than after stimulation (M = 23.45, SE = 1.01) in both tDCS sessions. Similarly, a main effect of time on PANAS Negative scores ($F(1,19) = 10.56$, $p < 0.01$, $\eta^2 = 0.070$) showed that scores were higher before (M = 11.28, SE = 0.28) than after tDCS in both the sham and the active tDCS condition (M = 10.53,

SE = 0.12). There were no further main or interaction effects of Time and tDCS condition for the PANAS Positive, PANAS Negative, or STAI-State scores (all p >0.05).

4.3.3. Study 2. Effects of offline tDCS on reward learning on the IBLT

4.3.3.1. Learning rate Win- and Loss-volatile blocks

As observed in the online tDCS study, there was a two-way interaction of Block x Outcome valence for learning rates ($F(1,19) = 37.12$, p <0.001, $\eta^2 = 0.158$). As predicted, win learning rates were higher in the Win-volatile blocks than in the Loss-volatile blocks ($t(79) = 5.25$, p <0.001), while loss learning rates were higher in the Loss-volatile blocks than in the Win-volatile blocks ($t(79) = 6.14$, p <0.001). Moreover, win learning rates were higher than loss learning rates in the Win-volatile blocks ($t(79) = 7.73$, p <0.001), and vice versa in the Loss-volatile blocks ($t(79) = 3.56$, p <0.001).

There was a main effect of tDCS condition ($F(1,18) = 5.06$, p <0.05, $\eta^2 = 0.028$), reflecting overall lower learning rates in the active tDCS session compared to the sham tDCS session (see Figure 4.4.). However, if win-loss baseline learning rates were included as a covariate, this main effect was no longer significant ($F(1,18) = 2.03$, $p = 0.172$), indicating that differences in learning rates prior to sham and active stimulation were responsible for the apparent effect of tDCS condition. In contrast with the online tDCS analyses, there were no interaction effects of tDCS on learning rates (all p >0.05). There was therefore no evidence for an effect of offline tDCS on IBLT performance.

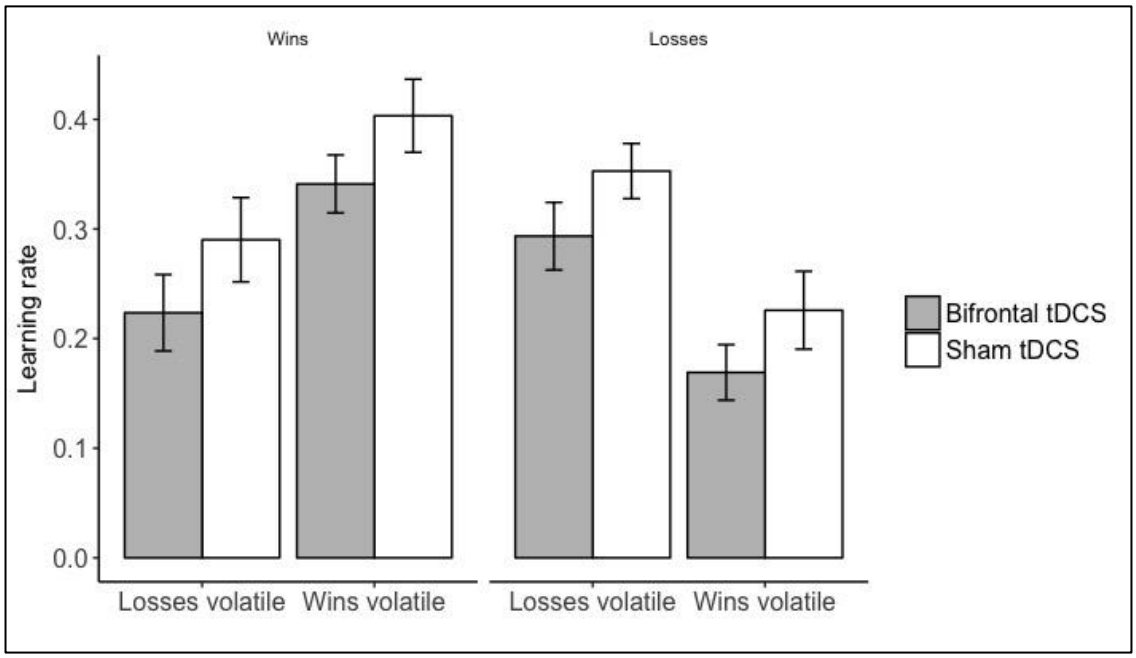

Figure 4.4. Win and loss learning rates in the Win- and Loss-volatile blocks by tDCS condition for the offline tDCS study. Error bars represent the standard error of the mean.

4.3.3.2. Learning rate Both-volatile blocks

For the Both-volatile blocks, there was a main effect of Time ($F(1,18) = 7.18$, $p < 0.05$, $\eta^2 = 0.053$), with learning rates being higher before (M = 0.41, SE = 0.03) than after tDCS (M = 0.32, SE = 0.03). There was a trend towards a main effect of tDCS condition ($F(1,18) = 3.47$, $p = 0.079$), suggesting that before and after stimulation learning rates tended to be lower in the active tDCS session than the sham tDCS session (see Figure 4.5.). This supports the hypothesis that lower learning rates in the Win- and Loss-volatile blocks with bifrontal compared to sham tDCS reflected an overall chance difference in learning rates between the two sessions rather than an effect of stimulation. There were no further main effects or interaction effects of tDCS condition, Outcome valence, and Time (all $p > 0.05$).

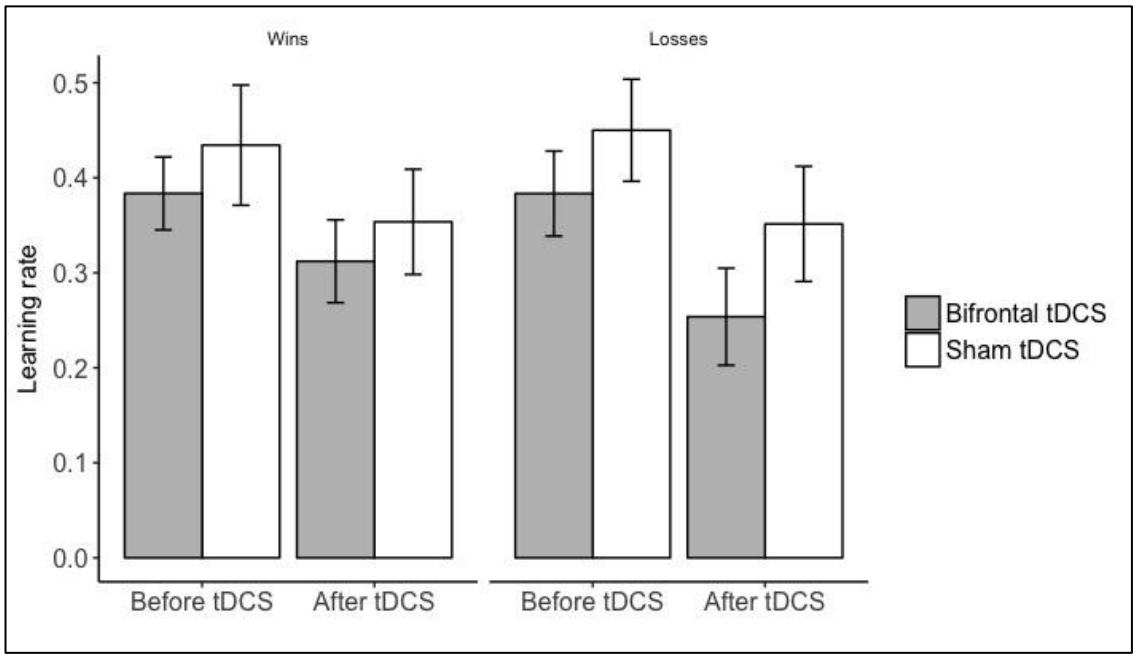

Figure 4.5. Win and loss learning rates in the Both-volatile blocks by tDCS condition. Error bars represent the standard error of the mean.

4.3.3.3. Choice stochasticity

In addition to the learning rates, I examined the influence of offline tDCS on inverse temperature values across task blocks. Reproducing results from the online tDCS group, there were no significant main effects of tDCS on the inverse temperature or any interactions of tDCS, Block, and Time (all p <0.05).

4.3.3.4. Win-driven choice behaviour and total winnings

A main effect of Block was again found for win-driven choice behaviour on the IBLT (F(1,19) = 139.05, p <0.001, η^2 = 0.775), reflecting a greater proportion of win-driven choices in the Win-volatile blocks (M = 0.70, SE = 0.01) than in the Loss-volatile blocks (M = 0.34, SE = 0.01). There was a trend towards a main effect of tDCS condition (F(1,19) = 3.65, p = 0.071), suggesting that the proportion of win-driven choices tended to be greater with active tDCS (M = 0.53, SE = 0.02) than with sham tDCS (M = 0.52, SE = 0.02). There was no evidence for an interaction of tDCS condition x Block (F(1,19) = 1.76, p = 0.200) in the Win-volatile and Loss-volatile blocks. There were no further main or

interaction effects in the Win-volatile, Loss-volatile, or Both-volatile blocks (all $p > 0.05$). In addition, while participants tended to win more money in the active tDCS than the sham tDCS condition, this difference was not significant ($t(19) = 1.72$, $p = 0.101$).

4.3.3.5. Mood and anxiety measures

Next, I investigated acute changes in mood and anxiety associated with tDCS condition. For PANAS Positive scores, there was a significant effect of Time ($F(1,19) = 16.23$, $p < 0.001$, $\eta^2 = 0.061$), with scores being higher before (M = 28.40, SE = 0.99) than after tDCS (M = 25.20, SE = 0.98). For the PANAS Negative scores there was a main effect of Time ($F(1,19) = 12.39$, $p < 0.01$, $\eta^2 = 0.020$) as well as an interaction of tDCS condition and Time ($F(1,19) = 5.72$, $p < 0.01$, $\eta^2 = 0.013$). Paired *t*-tests showed that there was a significant change in PANAS Negative scores over time with sham tDCS ($t(19) = 3.39$, $p < 0.01$), with scores being higher before (M = 12.95, SE = 0.66) than after tDCS (M = 11.55, SE = 0.57). In contrast, there were no significant differences of PANAS Negative scores before and after bifrontal tDCS ($t(19) = 0.38$, $p = 0.705$). There thus appeared to be a greater decrease in negative affect with sham tDCS than with active tDCS. However, it should be noted that there was a marginally significant difference between PANAS Negative scores in the sham and active tDCS session prior to stimulation ($t(19) = 2.01$, $p = 0.059$), with scores being lower before bifrontal tDCS (M = 12.00, SE = 0.60) than before sham tDCS (M = 12.95, SE = 0.66). There may therefore have been greater scope for a decrease in scores in the sham session than in the active tDCS session. There were no significant main or interaction effects of Time and tDCS condition on STAI-State scores (all $p > 0.05$).

4.3.3.6. Comparison of online and offline tDCS effects on learning rates

The significant effect of increased learning rates observed with online but not offline tDCS suggests that the effect of stimulation depends on the cognitive state of the participant

during tDCS. To formally test this hypothesis, I contrasted the datasets from the online and offline tDCS conditions directly. A repeated-measures ANOVA with win learning rates in the Loss-volatile block as the outcome variable, Group (online tDCS versus offline tDCS) and tDCS order as between-subject predictors, and tDCS condition and Run as within-subject predictors was carried out. There was a Group x tDCS condition x Time interaction ($F(1,36) = 5.34$, $p < 0.05$, $\eta^2 = 0.004$). Follow-up analyses showed that the Group x tDCS condition was significant in Run 2 ($F(1,36) = 10.89$, $p < 0.01$, $\eta^2 = 0.037$) but not Run 1 ($F(1,36) = 1.22$, $p = 0.277$). Specifically, learning rates were higher with active than sham tDCS in Run 2 for the online tDCS group ($t(79) = 3.03$, $p < 0.01$), whereas the opposite pattern was found in the offline tDCS group ($t(79) = -3.35$, $p < 0.01$). In addition, there was a significant Group x tDCS condition x Block interaction ($F(1,36) = 8.07$, $p < 0.001$, $\eta^2 = 0.006$) and a Group x tDCS condition x Outcome valence interaction ($F(1,36) = 3.01$, $p < 0.05$, $\eta^2 = 0.003$). As expected, the results indicated that the interactions of tDCS condition and Block or Outcome valence were significant in the online tDCS group but not the offline tDCS group. These findings support the hypothesis that bifrontal tDCS results in increased reward learning rates when applied during cognitive engagement but not when administered at rest.

4.3.4. Study 3. Replication of online tDCS effects

It was thus demonstrated that online but not offline tDCS could increase learning rates for stable win outcomes when applied to bifrontal DLPFC. In Study 3, it was attempted to replicate this finding in a new sample of healthy adults.

4.3.4.1. Participant characteristics

As shown in Table 4.2., depression and trait anxiety scores of the Study 3 sample did not differ significantly from the Study 1 sample. In the replication study, participants had higher PANAS Positive scores on average before the start of the sham tDCS session than

participants in Study 1. There were no significant differences in either PANAS Negative or STAI-State scores between the study samples.

Similar to Study 1, participants had a higher learning rate for win outcomes than loss outcomes in the Both-volatile block at baseline in the first session ($t(19) = 2.66$, $p < 0.05$), indicating that participants had an initial bias towards learning from positive compared to negative outcomes. These baseline learning rates did not differ from Study 1 for either wins ($t(37) = 0.72$, $p = 0.476$) or losses ($t(38) = 1.76$, $p = 0.086$). Therefore, any differences in the results between Study 1 and Study 3 were unlikely to be caused by differences in baseline learning performance.

Finally, baseline learning rates were compared for the sham and the active tDCS session in the replication sample to assess potential differences in learning prior to the two stimulation conditions. Before stimulation, there was an interaction of Outcome valence and tDCS condition ($F(1,19) = 9.61$, $p < 0.01$, $\eta^2 = 0.056$) for the learning rates in the replication study. Paired *t*-tests indicated that loss learning rates were significantly higher before the sham tDCS session than before the bifrontal tDCS session ($t(19) = 3.83$, $p < 0.01$), whereas there was no difference for win learning rates ($t(19) = -0.11$, $p = 0.916$). In addition, the learning rates for win outcomes were higher than for losses before stimulation in the bifrontal tDCS sessions ($t(19) = 3.67$, $p < 0.01$), whereas there was no difference between win and loss learning rates in the sham tDCS condition ($t(19) = 0 = 0.71$, $p = 0.485$). These findings suggest that participants may have presented with a greater tendency to learn from positive relative to negative outcomes before the active tDCS session than the sham tDCS session.

Table 4.2. Mean (SE) baseline trait and state anxiety and mood questionnaire scores for participants in Study 1 and Study 3. *t*-values and *p*-values represent the difference in scores between the two studies.

Questionnaire	Study 1	Study 3	*t*-value	*p*-value
BDI	5.11 (1.20)	3.95 (0.79)	0.43	0.67
STAI-Trait	37.47 (2.05)	34.95 (1.30)	0.19	0.85
PANAS Positive				
Sham tDCS session	26.85 (1.09)	31.32 (1.59)	-2.15	0.04
Bifrontal tDCS session	27.20 (1.19)	30.59 (1.53)	-1.71	0.10
PANAS Negative				
Sham tDCS session	10.90 (0.25)	11.63 (0.47)	-1.15	0.26
Bifrontal tDCS session	11.65 (0.50)	11.68 (0.43)	-0.22	0.83
STAI-State				
Sham tDCS session	29.95 (1.04)	28.73 (1.41)	0.47	0.64
Bifrontal tDCS session	30.00 (1.38)	29.77 (1.68)	-0.22	0.83

4.3.4.2. Learning rates Win- and Loss-volatile blocks

First, I investigated the influence of tDCS condition on the learning rates across task blocks and sessions. There was a main effect of Outcome valence ($F(1,18) = 4.25$, $p < 0.05$, $\eta^2 = 0.027$), with participants having higher learning rates for win outcomes (M = 0.27, SE = 0.02) than for loss outcomes (M = 0.22, SE = 0.01). Again, an interaction of Block x Outcome valence was observed ($F(1,18) = 68.85$, $p < 0.001$, $\eta^2 = 0.271$), with win learning rates being higher in Win-volatile blocks than Loss-volatile blocks ($t(79) = 7.27$, $p < 0.001$) and loss learning rates being higher in the Loss-volatile blocks than in the Win-volatile blocks ($t(79) = 9.15$, $p < 0.001$). Moreover, win learning rates were higher than loss learning rates in the Win-volatile blocks ($t(79) = 10.22$, $p < 0.001$), while loss learning rates were higher than win learning rates in the loss-volatile blocks ($t(79) = 6.16$, $p < 0.001$). Similar to Studies 1 and 2, this demonstrated that participants tended to learn faster from volatile than stable stimuli. In contrast to Study 1, there was no significant interaction between tDCS condition and Block ($F(1,18) = 0.00$, $p = 0.930$) or tDCS condition and

Outcome valence ($F(1,18) = 1.30$, $p = 0.269$).

A main effect of Run was found ($F(1,18) = 8.67$, $p < 0.01$, $\eta^2 = 0.006$), with learning rates being higher in the first two blocks (M = 0.25, SE = 0.02) than in the second two blocks (M = 0.23, SE = 0.02). Analyses were therefore again run separately for both Runs of the IBLT. In Run 1, there was a main effect of tDCS condition ($F(1,18) = 4.83$, $p < 0.05$, $\eta^2 = 0.021$), with learning rates being higher in the active tDCS session (M = 0.27, SE = 0.02) than the sham tDCS session (M = 0.23, SE = 0.02). However, this effect was not maintained if baseline win-loss learning rates were included as a covariate ($F(1,18) = 0.14$, $p = 0.708$). There were thus no indications that active tDCS modified learning rates during stimulation (see Figure 4.6.A.). In Run 2, there was no main effect of tDCS condition ($F(1,18) = 0.01$, $p = 0.939$), nor an interaction effect of tDCS condition x Outcome valence ($F(1,18) = 0.26$, $p = 0.618$). The effects of active tDCS on reward learning after stimulation demonstrated in Study 1 were thus not replicated in Study 3 (see Figure 4.6.B.).

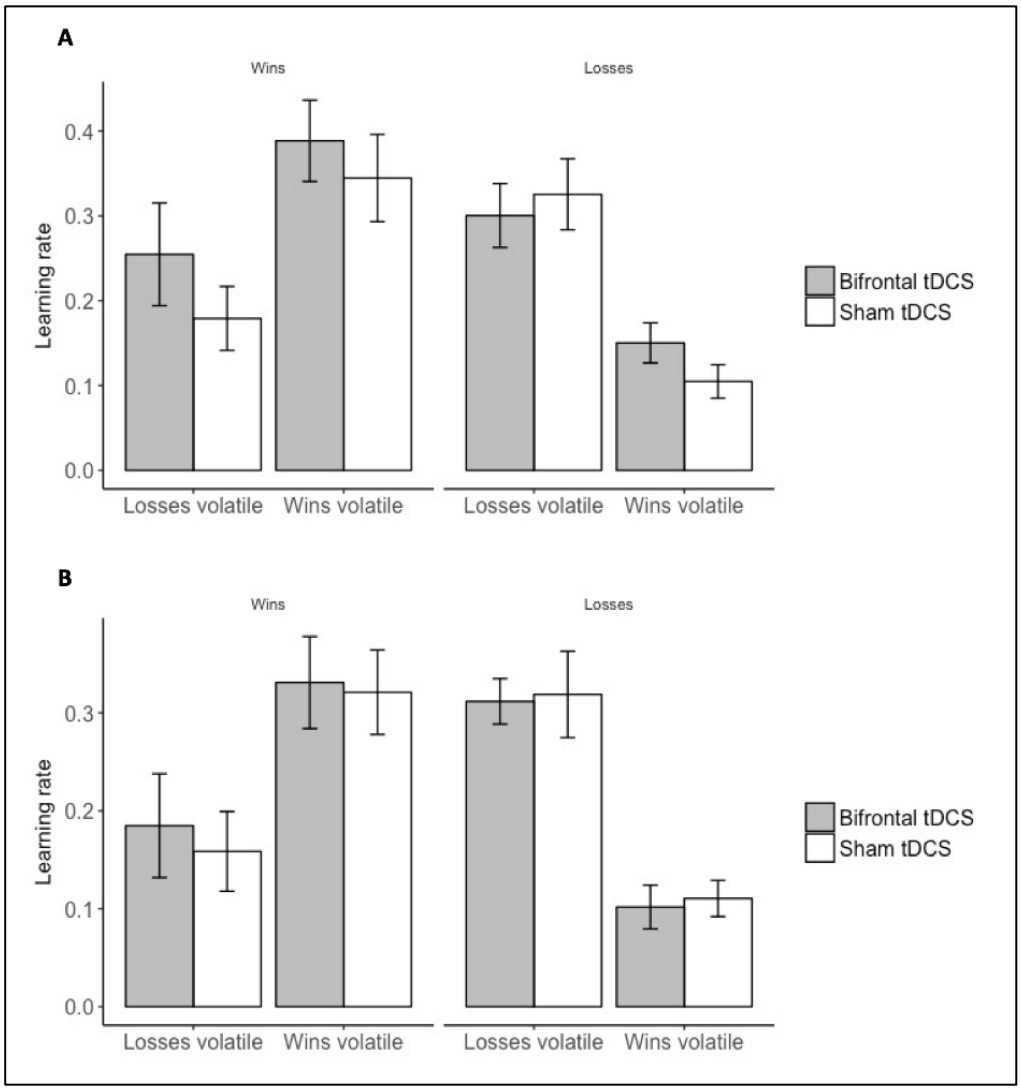

Figure 4.6. Effects of online tDCS on in win and loss learning rates by tDCS condition and Block for the Replication study (Study 3) in the Win- and Loss-volatile blocks. **(A)** Learning rates during tDCS (Run 1). **(B)** Learning rates after tDCS (Run 2). Error bars represent the standard error of the mean.

4.3.4.3. Learning rates Both-volatile blocks

For the Both-volatile blocks, there was a main effects of Time ($F(1,18) = 11.62$, $p < 0.01$, $\eta^2 = 6.46$), with learning rates being higher before (M = 0.33, SE = 0.03) than after tDCS (M = 0.24, SE = 0.03). In addition, a main effect of Outcome valence was observed ($F(1,18) = 13.14$, $p < 0.01$, $\eta^2 = 0.973$), with higher learning rates for wins (M = 0.35, SE = 0.03) than losses (M = 0.22, SE = 0.02). Finally, there was a tDCS condition x Time x

Outcome valence interaction ($F(1,18) = 5.97$, $p < 0.05$, $\eta^2 = 0.021$). Further analyses indicated that there was a significant tDCS x Time interaction for loss outcomes ($F(1,18) = 10.43$, $p < 0.01$, $\eta^2 = 0.077$) but not win outcomes ($F(1,18) = 0.35$, $p = 0.563$). Paired *t*-tests showed that there was a decrease in loss learning rates with sham tDCS ($t(19) = -2.71$, $p < 0.05$), whereas no change over time was observed with active tDCS ($t(19) = 0.92$, $p = 0.371$). In addition, learning rates for losses were higher before sham tDCS than before bifrontal tDCS ($t(19) = 3.83$, $p < 0.01$), whereas there were no differences between the two conditions after stimulation ($t(19) = -0.01$, $p = 0.990$). This greater reduction in learning from negative outcomes with sham tDCS than active tDCS was therefore most likely due to higher baseline loss learning rates prior to stimulation, thereby providing greater scope for a decrease in learning rates (see Figure 4.7.).

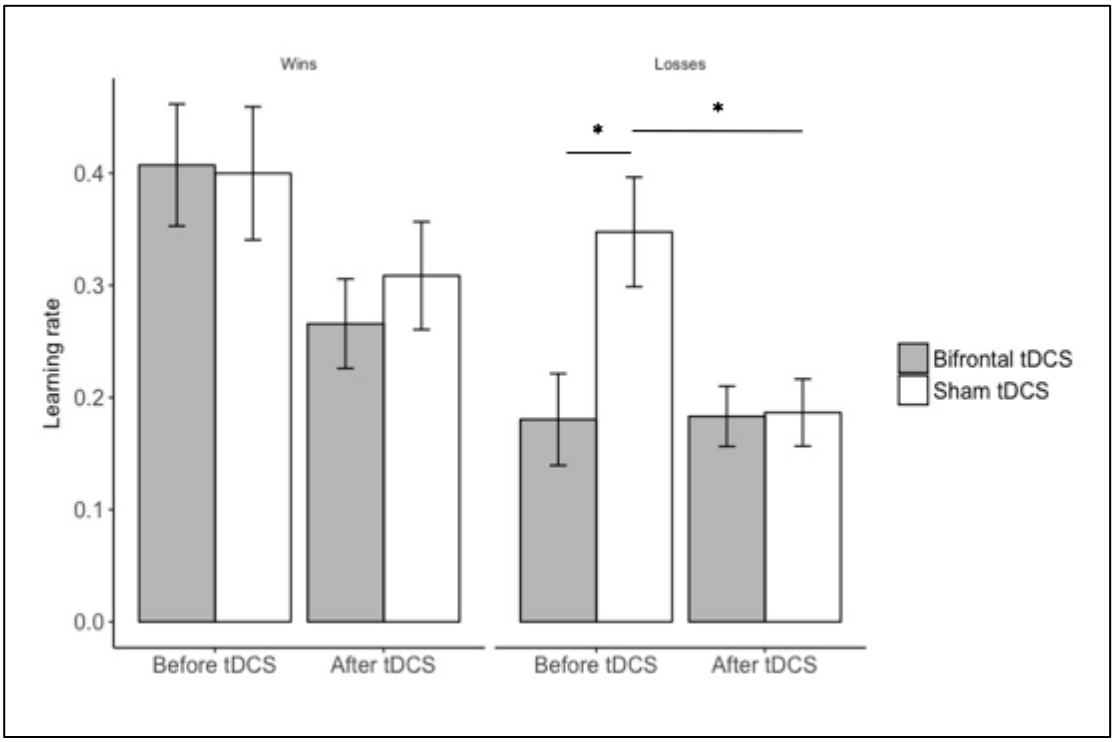

Figure 4.7. Learning rates in the Both-volatile blocks by tDCS condition, Outcome valence and Time in the Replication study. Error bars represent the standard error of the mean (* $p < 0.01$).

4.3.4.4. Choice stochasticity

Similar to Study 1, no main effects of tDCS were observed for choice stochasticity as indicated by the inverse temperature values ($F(1,18) = 0.15$, $p = 0.705$). In addition, there were no significant interactions of tDCS condition, Block, and/or Time on the inverse temperature values (all $p > 0.05$).

4.3.4.5. Win-driven choice behaviour and total winnings

For the proportion of win-driven choices following trials in which the win and loss outcomes were associated with the same shape, there was a main effect of Block ($F(1,18) = 391.98$, $p < 0.001$, $\eta^2 = 0.833$). Participants tended to make more win-driven choices in the Win-volatile blocks (M = 0.72, SE = 0.01) than in the Loss-volatile blocks (M = 0.32, SE = 0.01). There were no significant main effects of tDCS ($F(1,18) = 0.25$, $p = 0.626$) or interaction effects of tDCS condition x Block ($F(1,18) = 0.43$, $p = 0.519$). There was therefore no replication of the increase in win-driven choice behaviour in the Loss-volatile blocks observed in Study 1. In the Both-volatile blocks, there were no main or interaction effects of tDCS condition or Time (all $p > 0.05$). There was no significant difference in the total amount of money won in the active and sham tDCS sessions ($t(19) = 0.34$, $p = 0.730$).

4.3.4.6. Mood and anxiety measures

Within the replication study, there was a decline in PANAS Positive scores over time ($F(1,19) = 15.51$, $p < 0.001$, $\eta^2 = 0.050$) in both sessions. There were no further changes over time or any main or interaction effects of tDCS on the PANAS Positive, PANAS Negative, or STAI-State scores (all $p > 0.05$).

4.3.5. Pooled data for the effects of online tDCS on learning rates

4.3.5.1. Win-volatile and Loss-volatile blocks

Finally, I studied the main effects of online tDCS on learning rates by pooling the data from Study 1 and Study 3 in one analysis to increase the sample size and thereby improve

power to estimate the effects of stimulation. There was a main effect of Outcome valence (M = 6.06, p <0.05, η^2 = 0.016), with participants having higher learning rates for win outcomes (M = 0.30, SE = 0.01) than for loss outcomes (M = 0.25, SE = 0.01). Consistent with previous analyses, a Block x Outcome valence was observed (F(1,38) = 27.55, p <0.001, η^2 = 0.028), indicating that learning rates were higher for volatile stimuli than for stable stimuli across the blocks.

Crucially, there was a tDCS condition x Outcome valence interaction (F(1,38) = 5.47, p <0.05, η^2 = 0.004). Paired t-tests indicated that win learning rates were significantly higher with active tDCS than sham tDCS (t(159) = 2.94, p <0.01), whereas there was no difference for loss learning rates (t(159) = 0.588, p = 0.558). In addition, win learning rates were higher than loss learning rates with active tDCS (t(159) = 3.08, p <0.01) but not sham tDCS (t(159) = 1.14, p = 0.256). While the tDCS condition x Outcome valence x Block interaction was not significant (F(1,38) = 1.76, p = 0.192), the effect appeared to be strongest in the Loss-volatile blocks (see Figure 4.8.). The main interaction of tDCS condition and Outcome valence remained significant when baseline win-loss learning rate was included as a covariate (F(1,38) = 5.46, p <0.05, η^2 = 0.004). When Study (Study 1 versus Study 3) was included as a between-subjects variable, the Study x tDCS condition x Outcome valence interaction was not significant (F(1,36) = 0.612, p = 0.439), suggesting that the effect of tDCS on win learning rates was not solely driven by the results obtained from Study 1.

Again, there was a main effect of Run (F(1,38) = 5.02, p <0.05, η^2 = 0.01), with learning rates being higher in the first three blocks (M = 0.29, SE = 0.01) than in the last three blocks (M = 0.27, SE = 0.01). As in the prior analyses, effects of tDCS were therefore explored separately for the IBLT blocks carried out during and after tDCS. When the runs

were investigated in isolation, there was no interaction effect of tDCS condition x Outcome valence in Run 1 ($F(1,38) = 2.81$, $p = 0.102$) or Run 2 ($F(1,38) = 3.33$, $p = 0.762$).

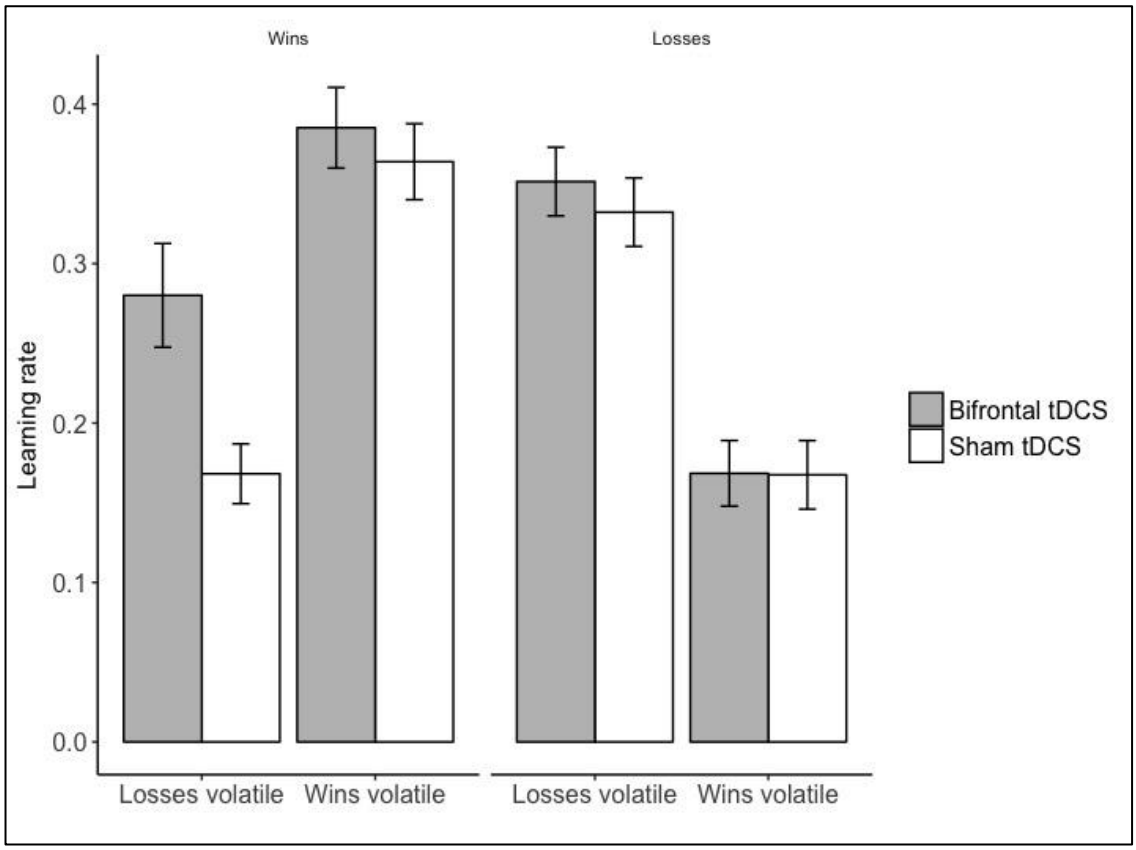

Figure 4.8. Learning rates in the Win- and Loss-volatile blocks for Study 1 and 3 combined. Error bars represent the standard error of the mean.

4.3.5.2. Both-volatile blocks

For the Both-volatile blocks, there was a main effect of Outcome valence ($F(1,38) = 16.08$, $p < 0.001$, $\eta^2 = 0.069$), with higher learning rates for wins (M = 0.35, SE = 0.03) than for losses (M = 0.22, SE = 0.02). In addition, there was a main effect of Time ($F(1,38) = 15.55$, $p < 0.001$, $\eta^2 = 0.034$), such that learning rates were higher before (M = 0.33, SE = 0.03) than after tDCS (M = 0.24, SE = 0.02). Finally, there was a tDCS condition x Outcome valence interaction ($F(1,38) = 8.92$, $p < 0.01$, $\eta^2 = 0.012$). It appeared that loss learning rates were higher in the sham tDCS condition than in the bifrontal tDCS condition

($t(79) = 2.30$, $p < 0.05$), whereas there were no differences in win learning rates ($t(79) = -1.13$, $p = 0.262$). Win learning rates were higher than loss learning rates in both the bifrontal tDCS condition ($t(79) = 5.51$, $p < 0.001$) and the sham tDCS condition ($t(79) = 2.25$, $p < 0.05$). However, there were no interactions with Time, suggesting that there was no change from pre- to post-stimulation learning rate as a function of tDCS (see Figure 4.9.).

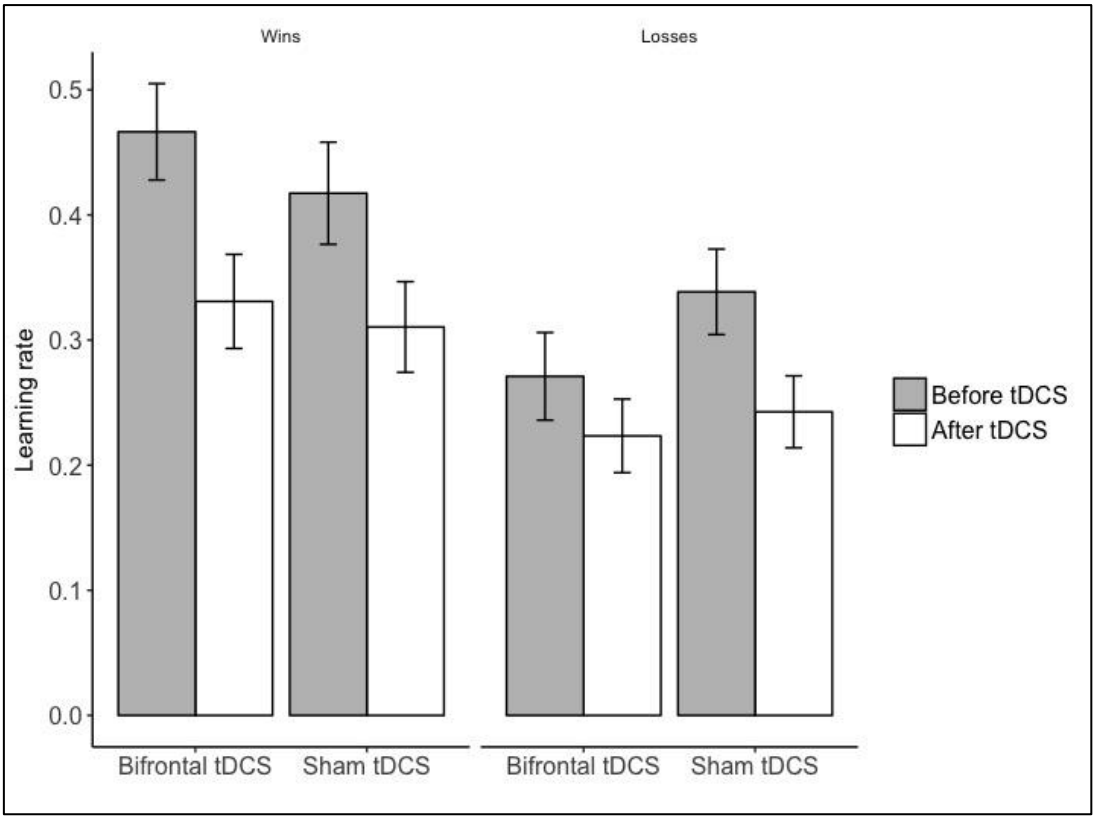

Figure 4.9. Learning rates in the Both-volatile blocks before and after sham or bifrontal tDCS. Error bars represent the standard error of the mean.

4.3.6. Power calculations

In order to determine the minimum detectable effect size in the present studies, power calculations were carried out using G*Power 3.1. software (Faul et al., 2007). For a two-tailed paired comparison using an alpha of 0.05 and power of 0.80, it was found that an effect size of 0.660 would be required to detect a significant effect in a sample of 20 participants. For a sample size of 40 participants, as in the pooled dataset described above,

the required effect size was reduced to 0.454. Finally, an effect size of 0.909 would be needed in order to detect a between-subjects effect when contrasting the different groups with 20 participants in each sample.

4.3.7. Participant blinding efficacy across studies

Finally, as it has previously been suggested that participant blinding may be inadequate when applying tDCS at 2 mA (O'Connell et al., 2012), it was investigated whether participants could correctly identify the stimulation condition across sessions. In the online tDCS study (Study 1), it was found that participants identified the true active tDCS condition correctly in 90% of all sessions, whereas the sham tDCS condition was identified in 60% of the sessions. Seven participants within this study indicated that they believed they had received active tDCS condition in both sessions, although all participants had been informed that they would be given active tDCS in one session and sham tDCS in the other. This is likely to have inflated the percentage of correct guesses in the active tDCS condition. In the offline tDCS study (Study 2), a lower number of participants was able to identify the tDCS condition, with volunteers guessing the active tDCS condition correctly in 77% and the sham tDCS condition in 54% of all sessions. Here, 6 out of 20 participants chose the same condition in both sessions. In the replication study (Study 3), blinding efficacy appeared to be slightly greater than in the other two experiments, with participants guessing the correct condition in 68% of the active tDCS sessions and 59% of the sham tDCS sessions and 10 participants choosing the same tDCS condition across the two sessions.

4.4. Discussion

4.4.1. Main findings

This study aimed to test the hypothesis that non-invasive brain stimulation of the dorsolateral prefrontal cortex and connected regions could alter learning rates in healthy

adults. The findings from Study 1 indicated that online bifrontal stimulation of the DLPFC during a probabilistic learning task may increase learning rates for uninformative and informative positive outcomes. The specificity of the changes in learning rates for reward but not loss outcomes associated with tDCS suggest that these are not due to either general changes in cognitive control or working memory processes or distraction from the task related to stronger physical sensations associated with active tDCS. In a replication study applying the same tDCS parameters, however, no significant effects of non-invasive brain stimulation on reward learning rates were found. Previously, it has been suggested that a failure to replicate can be due to the absence of a true effect or the result of low statistical power in single replication studies (Maxwell, Lau, & Howard, 2015). As described in Section 4.3.6., a moderate to large effect size would be needed to detect a significant effect in a sample of 20 participants. It is therefore possible that the lack of replication is due to this study being underpowered to detect the small effect of tDCS on learning rates as observed in Study 1. When pooling the two datasets, the increase in learning rates for win outcomes with active compared to sham tDCS again became apparent. Furthermore, it was found that this effect did not interact with the specific study investigated (i.e., the initial study versus the replication study), suggesting that an increase in win learning rates was present in both studies rather than being driven solely by Study 1. Nevertheless, given the lack of replication in a study sample independent from Study 1, the findings of a change in IBLT performance with tDCS reported here should not be overinterpreted. In contrast to the findings from the online stimulation studies, tDCS applied at rest was ineffective in modifying learning parameters. In line with previous reports, these results suggest that tDCS combined with cognitive engagement may be a superior method for altering cognitive skills compared to offline tDCS (Martin et al., 2014).

Investigations of win-driven choice behaviour in Study 1 confirmed the enhancing effects of online tDCS on learning from positive outcomes, although this change appeared to be limited to Loss-volatile blocks. Although this finding thus did not replicate global increases in learning rates from win outcomes in the computational analyses, it should be noted that the effect of tDCS on reward learning rates did appear to be numerically greatest in the Loss-volatile blocks. Potentially, this may be due to a ceiling effect of win learning in blocks where win outcomes are highly informative, thereby restricting larger increases in learning in Win-volatile blocks. As win-driven choice behaviour is a less sensitive measure than learning rates, it is possible that analyses of this outcome were unable to detect the subtle changes in behaviour on the Win-volatile blocks. It should be noted that, as the number of win-driven and loss-driven choices are mirrored, this finding might also be interpreted as a decrease in loss-driven choices. Given that there was an increase in win learning rates but no alteration of loss learning rates, however, it is likely that these changes reflect a modification of win-driven rather than loss-driven behaviour. No alterations in randomness of behaviour were observed in any of the three tDCS studies as indicated by stable inverse temperature values across all blocks. This indicated that tDCS was able to specifically alter learning rates while leaving other computational parameters unchanged.

This is the first study to provide evidence that the effect of bifrontal tDCS on learning rates was specific to positive events, while learning from negative events was relatively unaffected. Previously, research has shown that positive and negative outcomes have different value functions, and correspondingly different underlying neural bases (Chen et al., 2015). For example, an ALE meta-analysis of 35 reinforcement studies indicated that appetitive and aversive prediction errors are coded through the brain by broadly segregated neural networks (Garrison, Erdeniz, & Done, 2013), although

additional research has indicated there are also overlapping regions for reward and punishment processing (Jessup & O'Doherty, 2014; Vickery, Chun & Lee, 2011). It can be speculated that the tDCS montage applied here mainly affected neural circuits involved in processing of positive rather than negative information. Neuroimaging studies may shed further light on the neural basis underlying the observed changes in reinforcement learning.

4.4.2. Study limitations

It should be noted that, in the current study, individuals tended to have higher learning rates for win outcomes than loss outcomes from the outset. It is currently unclear whether the findings from this research can be extrapolated to populations with different baseline learning rates. As discussed previously, it is possible that there is a different balance of learning rates for positive and negative outcomes in populations of interest, such as depressed individuals. Additional research is therefore needed to investigate whether initial biases towards processing positive or negative information interact with the application of tDCS.

In addition, although the first 10 trials of each block were excluded from the analyses to prevent possible inflation of learning rates at the start of each new task block, due to unfamiliarity with novel stimuli, it is possible that learning effects occurred for a more extended period within the first study session. That is, although participants had an opportunity to practice the IBLT prior to completing the full task, participants may have needed additional time in the first session to gain a clear understanding of the task. It is therefore possible that effects of tDCS on the Both-volatile blocks, which were carried out at the start and the end of the task, were obscured by differences in understanding of the task over time. In addition, results from Chapter 3 suggested that there may be an effect of fatigue, such that learning rates tend to decrease over time in the Both-volatile blocks.

Tiredness in the second run of the IBLT may therefore have negated potential alterations induced by tDCS.

Double-blinding was carried out in this study to account for potential placebo effects of the stimulation on task performance. Similar to previous reports of stimulation at 2 mA (O'Connell et al., 2012), however, it was found that a greater number of participants than would be expected by chance was able to correctly guess the tDCS conditions across the present three studies, indicating that participant blinding was suboptimal. It should be noted that participant blinding appeared to be more successful in Study 3 than in Study 1, which may have contributed to the failure of replication of increased reward learning rates with bifrontal tDCS. In other words, knowledge of the tDCS condition could hypothetically influence participants' behaviour on the task independent of the effects of the stimulation itself. However, participant unblinding does not explain the specificity of the effects of tDCS on reward but not loss learning rates. In addition, it is improbable that awareness of the tDCS condition would allow participants to consciously influence the complex computational parameters derived from the IBLT. In light of this specificity of the tDCS effects and the complexity of the task, it is therefore unlikely that the findings from Study 1 are solely due to ineffective blinding of the participants. In order to achieve greater blinding efficacy and increase confidence about the anatomical specificity of the effect of bifrontal tDCS on learning rates, however, future study designs may consider incorporating an additional active tDCS montage which does not affect excitability in the regions of interest.

Finally, it has previously been emphasised that faster learning does not necessarily correspond to improved learning and better performance (Scholl & Klein-Flügge, 2017). An example of this was found in the present study, with higher learning rates for wins being associated with a decrease in total winnings with bifrontal tDCS compared to sham

stimulation in Study 1. Thus, enhancing learning rates irrespective of the demands of the environment decreased performance within the constraints of the experimental task used here. The reported findings could therefore be interpreted as a failure to adapt win learning rates appropriately to the task conditions as a result of tDCS. In a broader context, however, increasing the estimated information content for positive events has the potential to ameliorate negative attention biases associated with psychiatric disorders. Additional research is clearly needed to determine whether the negative effects of tDCS on basic task performance observed in the present study would result in positive or negative outcomes in terms of attentional biases and general learning processes in both healthy adults and depressed individuals.

4.4.3. Conclusion

In this proof-of-concept study, the hypothesis that bifrontal stimulation of the DLPFC could alter learning rates for positive outcomes in healthy adults was partially supported. These encouraging findings demonstrate that reinforcement learning processes might be malleable with the application of non-invasive brain stimulation. Given the specificity of the effect observed here, it may be possible to shape reward learning in such a way as to enhance learning from positive information only. In the future, it would be of interest to investigate whether tDCS in its current form might be applied to change not only learning rates, but also consequent attentional biases associated with depression and anxiety.

Chapter 5. Effects of concurrent tDCS and IBLT training on reward learning and affective biases in healthy adults

5.1. Introduction

Cognitive models of depression emphasise the causal role of affective biases in the onset of disorders of depression (Beck, 1967). Negative biases are expressed as an increase in attention to and encoding of adverse information relative to positive information, which may result in low mood and related depressive symptoms (Dalgleish & Watts, 1990; Peckham, McHugh, & Otto, 2010). Due to their fundamental contribution to depressive symptomatology, there has been increasing interest in the development of novel interventions aiming to alter attentional biases. An influential approach has been the use of cognitive training paradigms to modify biased emotional processing directly. The most prominent example is Attentional Bias Modification (ABM), in which patients are asked to respond to a visual probe which is placed behind either a positive or a negative stimulus. By placing the probe consistently behind the positive stimulus, patients are taught to focus on positive stimuli in favour of negative stimuli, and should therefore develop a positive attention bias. Some successes have been reported in using this type of training to ameliorate symptoms of depression (Li et al., 2016; Yang et al., 2015; Yang et al., 2016) and to reduce the risk of recurrence of depressive episodes (Browning et al., 2012). However, several research studies have failed to find a positive effect of ABM tasks on depression (Baert et al., 2010; Beevers et al., 2015; De Voogd et al., 2016). In addition, two meta-analyses of ABM experiments have indicated that the effect size of training outcomes tends to be small (Hallion & Ruscio, 2011; Mogoase, David, & Koster, 2014). It has also been suggested that changes in attentional biases as a result of ABM may not generalise to untrained stimuli (Kruijt, Putman, & Van der Does, 2013), which would undermine its usefulness as an interventional strategy.

An alternative method to diminish the influence of attentional biases on mood and other symptoms of depression may be to focus on changing the underlying factors leading to the development of affective biases. As discussed in Chapter 4, negative biases may stem from a belief that adverse events have a higher information content compared with positive events, resulting in preferential processing of negative outcomes. In the present study, I aim to investigate whether a cognitive training task can be used to alter people's estimates of how informative positive and negative outcomes are, which in turn could influence attentional biases and consequently measures of mood. The training used in the current study is an adaptation of the Information Bias Learning Task (IBLT), a reinforcement learning task described in detail in Chapters 3 and 4. In this task, the information content of positive and negative outcomes is manipulated by varying the volatility of their associations with different visual stimuli, with higher changeability increasing the information content of individual events. As a result, participants tend to learn faster from volatile than stable outcomes on the IBLT. Increased learning rates for volatile outcomes thus likely reflect higher information content estimates for changeable stimuli compared to constant stimuli. The IBLT may therefore have the potential to enhance information content estimates and associated learning rates for either rewarding or aversive stimuli by increasing the volatility of positive or negative outcomes, respectively. Here, I will examine the effects of Win-volatile and Loss-volatile IBLT training on learning parameters and associated attentional biases in healthy adults.

The impact of cognitive training could potentially be enhanced further by stimulating brain regions involved in performance of the task. By using non-invasive brain stimulation, specifically transcranial direct current stimulation (tDCS), to try and promote brain plasticity and facilitate learning, this could result in augmented outcomes compared to training in isolation. Previously, it has been demonstrated that tDCS may have

synergistic effects on the performance of cognitive tasks and depressive symptoms. For instance, Clarke and colleagues (2014) showed that concurrent anodal tDCS of the left dorsolateral prefrontal cortex (DLPFC) and ABM training was associated with greater changes in attentional biases in the trained direction (either towards or away from threatening stimuli) compared to sham stimulation in healthy adults. Similarly, tDCS targeting the dorsolateral prefrontal cortex was found to boost the impact of cognitive control training (CCT) on symptoms of Major Depressive Disorder, particularly in older patients and participants demonstrating high task performance (Brunoni et al., 2014). Another study found that the effects of combined CCT and tDCS of the DLPFC resulted in maintained improvements of depressive symptoms post-intervention, whereas each individual intervention had no such effect (Segrave et al., 2014). In healthy adults, on the other hand, it has been reported that a combination of prefrontal tDCS and cognitive training enhanced task performance to a greater degree than the two interventions applied in isolation during stimulation, although there was no change in learning after tDCS had ended stimulation (Martin et al., 2013). There is thus evidence to suggest that tDCS of the DLPFC might enhance the impact of computerised cognitive training paradigms, although the duration and generalisability of the changes have varied across studies and populations.

In addition to testing for differential effects of Win-volatile and Loss-volatile training on behaviour, I will therefore also examine potential amplifying effects of bifrontal tDCS on IBLT training for learning from positive and negative outcomes. Furthermore, I will explore whether learning rates induced by IBLT training, combined with active or sham tDCS, transfers to changes in performance on the Facial Expression Recognition Task (FERT; Harmer et al., 2001), a task that measures emotional processing biases. The FERT has been used extensively to characterize the effects of pharmacotherapeutic interventions on affective biases in both healthy adults and people

with depression (Bhagwagar et al., 2004; Browning et al., 2007; Harmer et al., 2003; Tranter et al., 2009). In the FERT, participants are briefly presented with images of emotional facial expressions, and are required to indicate as quickly and accurately as possible which emotion was shown. Depression has been associated with several changes in FERT task performance. For example, patients with MDD tend to have difficulties in accurately identifying the various emotions presented (Anderson et al., 2011), are slower in responding to sad faces (Gollan et al., 2008), and interpret neutral faces as happy less often than healthy controls (Douglas & Porter, 2010; Gollan et al., 2008; Surguladze et al., 2004). It is therefore of particular interest to assess whether a combination of tDCS and IBLT training could influence participants' recognition of facial expressions of happiness and sadness. Furthermore, potential changes in FERT task performance might correlate with alterations in mood and anxiety, which have been shown to be related to attentional biases (Dalgleish & Watts, 1990).

In summary, the study reported here aimed to answer the following questions: (1) can a training version of the IBLT enhance learning from positive and/or negative outcomes in healthy adults?; (2) does tDCS targeting bilateral dorsolateral prefrontal cortex enhance the effects of reinforcement learning training towards positive or negative stimuli?; (3) do alterations in learning rates following tDCS and training transfer to affective biases in the recognition of positive or negative facial expressions?; and (4) do potential effects of training alone or training combined with brain stimulation change measures of momentary mood in healthy adults? It was expected that Win-volatile training would increase learning rates from rewarding outcomes, whereas Loss-volatile training would increase learning from loss outcomes. In a proof-of-concept study (see Chapter 4), it was suggested that prefrontal tDCS may specifically enhance learning rates for positive outcomes. This effect was most clearly evident in the Loss-volatile blocks, possibly

reflecting a ceiling effect in the Win-volatile blocks. Hence, while it was therefore hypothesised that non-invasive brain stimulation might increase win learning rates across the two IBLT training types, based on results from the previous study, this effect was predicted to be strongest in the Loss-volatile training condition. In addition, it was expected that training and/or stimulation-induced changes in reward and loss learning rates would be maintained after cessation of the stimulation and training. Furthermore, it was predicted that increases in reward learning rates would transfer to enhanced identification of happy facial expressions and reduced responsiveness to negative facial expressions on the FERT, whereas increases in loss learning rates were expected to have the opposite effect. Finally, potential effects of the training and tDCS protocol on momentary mood and anxiety were assessed. It was expected that Win-volatile training and/or active tDCS might enhance mood through increasing information processing of positive events. In contrast, Loss-volatile training may have negative effects on acute mood and anxiety due to increased processing of negative events. However, it should be noted that previous studies have indicated that single-session tDCS does not tend to affect mood in healthy adults (Morgan et al., 2014; Plazier et al., 2012; Vanderhasselt et al., 2013). In addition, it has been suggested that interaction with the social environment following alterations in emotional processing is required to induce changes in mood (Pringle et al., 2011). It was therefore expected that any changes in mood and anxiety due to the training and stimulation procedure would be small in the present sample of healthy volunteers. Findings from this study may be used to inform the development of interventions aiming to reduce attentional biases associated with mood disorders.

5.2. Methods and Materials

5.2.1. Participants

Twenty healthy adults were recruited from the community for each training condition to participate in the Win-volatile training study (8 female, mean age = 24.45, SD ± 5.06) and the Loss-volatile training paradigm (4 female, mean age = 23.05, SD ± 6.05). Time and travel expenses were reimbursed at a rate of £10 per hour. Furthermore, participants received additional money through performing the Information Bias Learning Task. Study procedures received approval from the local ethics committee (CUREC; R48995/RE003).

5.2.2. General procedure

All participants took part in one sham and one bifrontal tDCS session, which were carried out at least one week apart. The order of tDCS conditions was counterbalanced and double-blinded in both of the training groups. Following provision of written informed consent, the researcher completed a tDCS safety screening form with the participant at the start of each session to rule out potential contraindications to non-invasive brain stimulation. Participants were then asked to complete Beck's Depression Inventory II (BDI; Beck, Steer, & Brown, 1996) and the Trait subscale of the State-Trait Anxiety Inventory (STAI-Trait; Spielberger et al., 1983), followed by the State version of the State-Trait Anxiety Inventory (STAI-State; Spielberger et al., 1983) and the Positive and Negative Affect Scales (PANAS; Watson, Clark, & Tellegen, 1988). Next, the researcher explained the Information Bias Learning Task and set up the tDCS equipment as described in Chapter 4. The participant was then asked to complete five blocks of the IBLT (details of the task provided in Chapter 3), which consisted of a first Both-volatile block, three Win-volatile or Loss-volatile blocks for the Win-volatile and Loss-volatile training paradigms, respectively, and a second Both-volatile block (see Figure 5.1.). Sham or active tDCS was applied for 20 minutes as outlined in Chapter 4 during performance of the three training

blocks. If the participant finished the training blocks before the stimulation had ended, they were asked to sit quietly with their eyes open for the remaining time. The Both-volatile blocks were included to measure baseline learning rates and sustained changes in learning parameters following the training and tDCS application. Once the five IBLT blocks had been completed, participants were again asked to fill out the PANAS and STAI-State questionnaires to measure potential changes in acute mood and anxiety. Next, participants completed the Facial Expression Recognition Task described in section 5.2.3 to assess possible alterations in affective biases following the stimulation and IBLT training. All participants then filled out a third set of the PANAS and STAI-State questionnaires to measure mood and anxiety changes following performance of the emotionally charged FERT. Finally, participants were asked to indicate which tDCS condition they believed they had received in each session. In addition, to assess confidence in their choice, participants were required to rate their certainty about their guess on a scale of 0 – 100%, with a higher value indicating greater certainty. Furthermore, participants had to rate how distracted they were by the stimulation when carrying out the IBLT on a scale of 0 – 10 (see Appendix E). Low scores indicated greater distraction from the task due to tDCS.

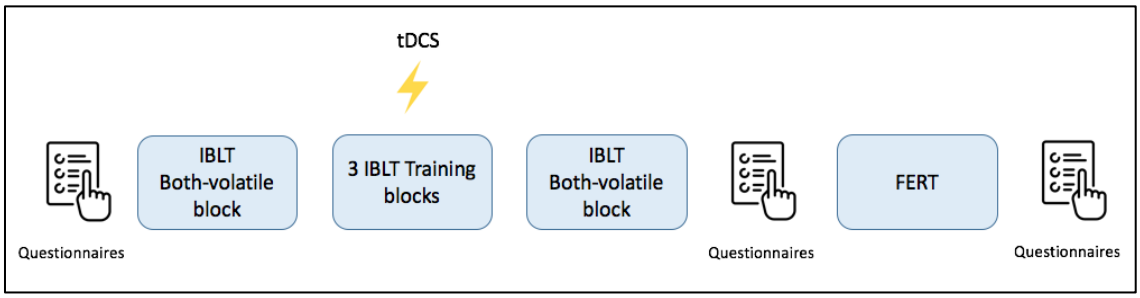

Figure 5.1. Schematic representation of the training procedure. All participants completed three training blocks of the Information Bias Learning Task (IBLT) while receiving bifrontal transcranial direct current stimulation (tDCS). Both-volatile blocks were presented preceding and following the training to assess near-transfer effects. Following the IBLT, participants performed the Facial Expression Recognition Task (FERT) as a measure of attentional bias. Participants were asked to fill out mood and anxiety questionnaires before the IBLT (Time 1), after the IBLT (Time 2), and after the FERT (Time 3).

5.2.3. Facial Expression Recognition Task (FERT)

In the Facial Expression Recognition Task, participants were asked to identify six basic emotions (anger, disgust, fear, happiness, sadness, and surprise) in facial expression stimuli derived from Ekman and Friesen's (1976) Pictures of Affect Series. All emotions were morphed between 0% intensity (neutral expression) and 100% intensity (full expression) in 10% increments. Each emotion was presented in 40 trials, with four examples of each emotion being shown for every intensity (10-100%). In addition, facial stimuli of neutral expressions (0% intensity) were presented in 10 trials, creating a total number of 250 trials. On every trial, the stimulus was presented for 500 ms, after which the image was replaced by a blank screen. Following stimulus presentation, participants were asked to indicate which emotion they believed was pictured by pressing a labelled key on a keyboard. Participants were instructed to make their choice as fast and accurately as possible. In its entirety, the task took approximately 10 minutes to complete. Measures derived from the task included accuracy (number of trials where the emotion was correctly

identified), reaction time (in milliseconds), and misclassifications (number of trials where the participant assigned an incorrect emotion to the presented stimulus) for every expression.

5.2.4. Statistical analyses

5.2.4.1. Learning parameters and win-driven choice behaviour

The effects of stimulation on learning rates and the inverse temperature values throughout the IBLT were examined using repeated-measures ANOVAs in R software (R Version 3.3.1.). Statistical analyses were initially carried out separately for the Win-volatile and Loss-volatile training groups, followed by analyses contrasting the data from the two training paradigms directly. Inverse logit transforms were applied to the learning rates, whereas a log transformation was used for the inverse temperature values.

To investigate changes in learning rates across the Training blocks for the Win-volatile and Loss-volatile training groups, Block (1, 2, or 3), Outcome valence (win outcome versus loss outcome), and tDCS condition (sham tDCS versus bifrontal tDCS) were applied as within-subject factors, while tDCS order (sham tDCS first versus bifrontal tDCS first) was included as a between-subject factor. For the Both-volatile blocks, the predictor variables were Time (pre- versus post-stimulation), Outcome valence, tDCS condition, and tDCS order. For the inverse temperature, independent variables in the Training blocks were Block, tDCS condition, and tDCS order, while the variables for the Both-volatile blocks were Time, tDCS condition, and tDCS order. In the pooled data analyses, Training type (Win-volatile versus Loss-volatile) was added as a between-subject variable to all the analyses. To account for the influence of individual differences in baseline task performance, all analyses were re-run with the inclusion of either baseline win learning rate, baseline loss learning rate, or baseline win – loss learning rate as obtained from the first Both-volatile block in the first testing session.

In addition to the computational analyses, win-driven behaviour as defined in Chapter 3 was investigated across the tDCS conditions. The proportion of win-driven choices was used as the outcome variable, with Block, tDCS condition, and tDCS order as the predictor variables.

5.2.4.2. Affective bias, mood, and anxiety measures

For the Facial Expression Recognition Task, repeated-measures ANOVAs were used to assess changes in Accuracy, Reaction time, and Misclassifications on the task. The independent variables were Emotion (Anger, Disgust, Fear, Happiness, Sadness, and Surprise), tDCS condition, and tDCS order. Alterations in acute mood and anxiety were investigated based on responses on the PANAS and STAI-State questionnaires. Effects of stimulation and training were examined by conducting repeated-measures ANOVAs with Time (1, 2, or 3), tDCS condition, and tDCS order as predictor variables for changes in PANAS Positive, PANAS Negative, and STAI-State scores.

5.3. Results

5.3.1. Baseline participant characteristics

To assess potential baseline differences in mood and anxiety prior to stimulation, scores on the BDI, STAI subscales, and PANAS questionnaires were compared between the Win-volatile and Loss-volatile training groups as well as the sham and active tDCS sessions. There were no significant differences between the two training groups on any of the questionnaires prior to the tDCS and training sessions (see Table 5.1.). In addition, there were no differences between baseline STAI-State or PANAS scores in the sham and active tDCS sessions within the two training groups (all $p > 0.05$).

To examine whether potential changes with Win-volatile versus Loss-volatile training might be influenced by baseline differences in learning abilities, baseline learning rates and inverse temperature parameters were compared for the two training types. In the

first IBLT block performed in the first session prior to stimulation, there were no differences in either win ($t(38) = 0.27$, $p = 0.791$) or loss learning rates ($t(36) = 0.73$, $p = 0.468$) between the Win-volatile and Loss-volatile training groups. However, there was a significant difference in baseline inverse temperature values ($t(26) = -2.67$, $p < 0.05$), with higher inverse temperature being found in the Win-volatile training group (M = 1.67, SE = 0.33) than in the Loss-volatile training group (M = 0.69, SE = 0.15). This indicates that choice behaviour of participants in the Loss-volatile training tended to be more stochastic than that of participants in the Win-volatile training prior to stimulation.

Table 5.1. Mean (SE) baseline questionnaire scores by Information Bias Learning Task training group. *t*-values and *p*-values represent differences between the Loss-volatile and Win-volatile training groups.

Questionnaire	Loss-volatile training	Win-volatile training	*t*-value	*p*-value
BDI-II	3.70 (1.33)	4.00 (33.15)	-0.19	0.853
STAI-Trait	34.35 (2.08)	33.15 (1.22)	0.50	0.622
PANAS Positive				
Sham tDCS session	33.40 (1.39)	32.35 (1.70)	0.48	0.636
Bifrontal tDCS session	33.45 (1.26)	33.60 (1.87)	-0.07	0.947
PANAS Negative				
Sham tDCS session	11.55 (0.62)	11.35 (0.43)	0.27	0.792
Bifrontal tDCS session	10.50 (0.18)	11.80 (0.63)	-1.98	0.060
STAI-State				
Sham tDCS session	27.40 (1.52)	26.15 (1.29)	0.63	0.535
Bifrontal tDCS session	26.35 (1.36)	26.60 (1.15)	-0.14	0.889

BDI-II; Beck's Depression Inventory II, PANAS; Positive and Negative Affect Scale; STAI, State-Trait Anxiety Inventory.

5.3.2. Win-volatile IBLT training

5.3.2.1. Learning rates

As expected, there was a main effect of Outcome valence on learning rates in the Training blocks ($F(1,18) = 135.39$, $p < 0.001$, $\eta^2 = 0.452$), with higher learning rates being observed

for win outcomes (M = 0.41, SE = 0.02) than for loss outcomes (M = 0.13, SE = 0.01). In addition, there was a main effect of Block ($F(2,36) = 3.85$, $p < 0.05$, $\eta^2 = 0.015$), such that learning rates were higher in Block 1 than in Block 3 ($t(79) = 2.46$, $p < 0.05$). There were no significant differences between Block 1 and Block 2 ($t(79) = 1.53$, $p = 0.131$) or Block 2 and Block 3 ($t(79) = 1.28$, $p = 0.203$). There appeared to be a trend towards a specific decrease in loss learning rates across the blocks with active bifrontal tDCS (see Figure 5.2.); however, the tDCS condition x Outcome valence x Block interaction was not significant ($F(1,18) = 2.02$, $p = 0.147$). There were no other main or interaction effects for Outcome valence, Block, or tDCS condition (all $p > 0.05$).

In the Both-volatile blocks, a main effect for Outcome valence was found ($F(1,18) = 15.15$, $p < 0.01$, $\eta^2 = 0.150$), with learning rates for win outcomes (M = 0.41, SE = 0.03) being higher than for losses (M = 0.26, SE = 0.02). In addition, there was a main effect of Time ($F(1,18) = 10.88$, $p < 0.01$, $\eta^2 = 0.045$), such that learning rates were higher before (M = 0.37, SE = 0.03) than after the training and tDCS (M = 0.30, SE = 0.03). There were no further significant main or interaction effects of tDCS condition, Time, Block, and/or Outcome valence on the learning rates. Results were unchanged if baseline win, loss, or win – loss learning rates were included as a between-subject covariate in any of the analyses.

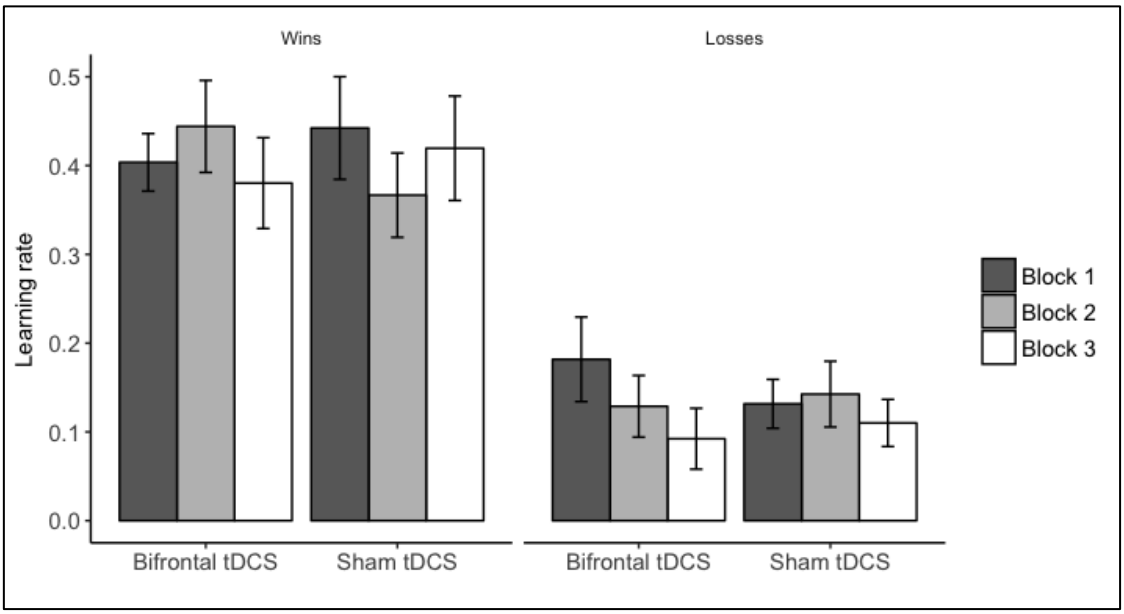

Figure 5.2. Learning rates in the Win-volatile training blocks by tDCS condition and outcome valence. Error bars represent the standard error of the mean.

5.3.2.2. Choice stochasticity

For the inverse temperature values, there were no significant effects of tDCS condition or Block in the Training blocks (all p >0.05). In the Both-volatile blocks there was a main effect of Time (F(1,18) = 10.47, p <0.01, η^2 = 0.071), with values being higher after (M = 16.76, SE = 2.56) than before training (M = 13.57, SE = 3.01). This suggests that choice randomness of the participants decreased with practice on the task. However, there was no significant effect of tDCS condition on the inverse temperature across the blocks (F(2,36) = 0.03, p = 0.972).

5.3.2.3. Win-driven choice behaviour and total winnings

For the proportion of win-driven choices, no significant main or interaction effects were found for either Block or tDCS condition in the Training blocks or the Both-volatile blocks (all p >0.05). Furthermore, there was no main effect of tDCS condition on the total amount of money won in the Training blocks (F(1,18) = 0.115, p = 0.738), nor an interaction effect of tDCS condition and Time in the Both-volatile blocks (F(1,18) = 2.45, p = 0.135). There

was thus no evidence that tDCS affected either computationally derived parameters or simple task performance in the Win-volatile training group.

5.3.2.4. Facial Expression Recognition Task

On the FERT, main effects of Emotion were found for Accuracy ($F(5,90) = 11.41$, $p < 0.001$, $\eta^2 = 0.266$), Reaction time ($F(5,90) = 18.29$, $p < 0.001$, $\eta^2 = 0.169$), and Misclassifications ($F(5,90) = 12.64$, $p < 0.001$, $\eta^2 = 0.327$). Overall, participants tended to be more accurate, faster, and less likely to misclassify emotions when presented with happy facial expressions compared to negative emotions (see Figure 5.3.). No further main effects or interactions were observed for tDCS condition and Emotion (all $p > 0.05$).

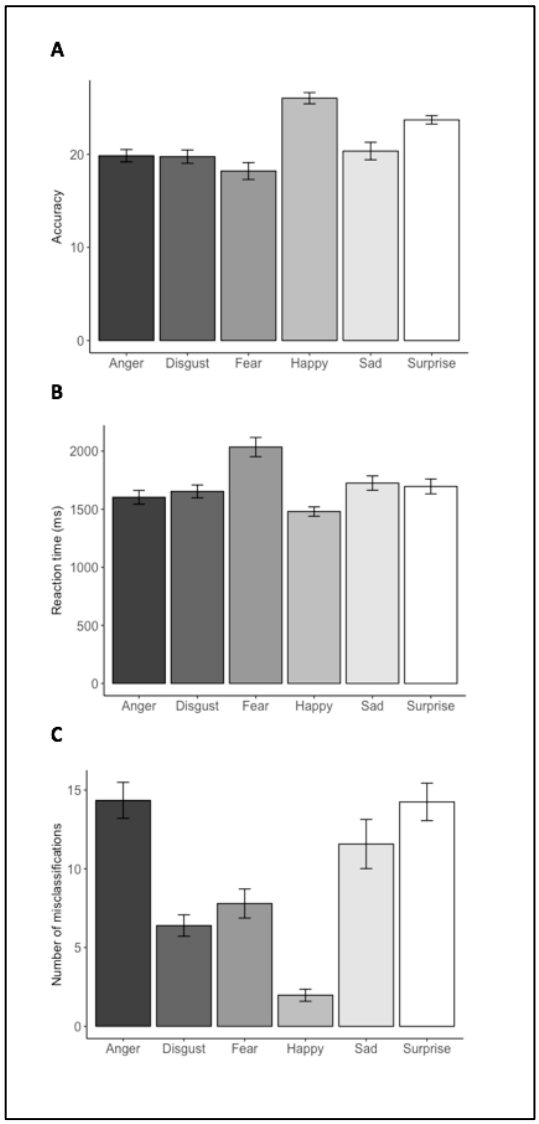

Figure 5.3. Performance on the Facial Emotion Recognition Task for all participants in the Win-volatile IBLT Training paradigm. **(A)** Number of accurate identifications by emotion. **(B)** Reaction time in milliseconds by emotion. **(C)** Number of misclassifications by emotion. Error bars represent standard error of the mean.

5.3.2.5. Mood and anxiety measures

There were no significant main or interaction effects of Time and tDCS condition on the PANAS Negative scores (all $p > 0.05$). However, there was a main effect of Time on both the PANAS Positive scores ($F(2,38) = 8.00$, $p < 0.01$, $\eta^2 = 0.024$) and the STAI-State scores ($F(2,38) = 3.40$, $p < 0.05$, $\eta^2 = 0.022$). Across the sessions, participants tended to have higher PANAS Positive scores at Time 1 (before the IBLT) than at Time 2 (after the IBLT) ($t(39) = 3.77$, $p < 0.01$) or Time 3 (after the FERT) ($t(39) = 3.87$, $p < 0.001$), whereas there was no difference between scores at Time 2 and 3 ($t(39) = -0.69$, $p = 0.497$). There was neither a main effect of tDCS condition ($F(1,19) = 0.49$, $p = 0.492$) nor a tDCS condition x Time interaction ($F(2,38) = 0.97$, $p = 0.389$) for the PANAS Positive scores. For the STAI-State scores, participants had a lower score at Time 1 than at Time 2 ($t(39) = -2.19$, $p < 0.05$) and Time 3 ($t(39) = -2.66$, $p < 0.05$), while there was again no difference between scores at Time 2 and Time 3 ($t(39) = -1.15$, $p = 0.259$). Again, there were no main or interaction effects of tDCS on participants' scores (all $p > 0.05$).

5.3.3. Loss-volatile IBLT training

5.3.3.1. Learning rates

For the learning rates in the Training blocks, there was a main effect of Outcome valence ($F(1,18) = 24.22$, $p < 0.001$, $\eta^2 = 0.112$), with learning rates being higher for loss outcomes (M = 0.32, SE = 0.01) than for win outcomes (M = 0.22, SE = 0.02). There were no other significant main or interaction effects of outcome valence, Block, and/or tDCS condition when investigating win or loss learning rates (see Figure 5.4.).

In the Both-volatile blocks, there was a main effect of Time ($F(1,18) = 5.77$, $p < 0.05$, $\eta^2 = 0.022$), as learning rates were higher before (M = 0.40, SE = 0.03) than after tDCS (M = 0.30, SE = 0.02). There were no further main or interaction effects on the task

for either win and loss learning rates (all p >0.05). Results were unchanged if the baseline win, loss, or win – loss learning rates were included as a covariate in the analyses.

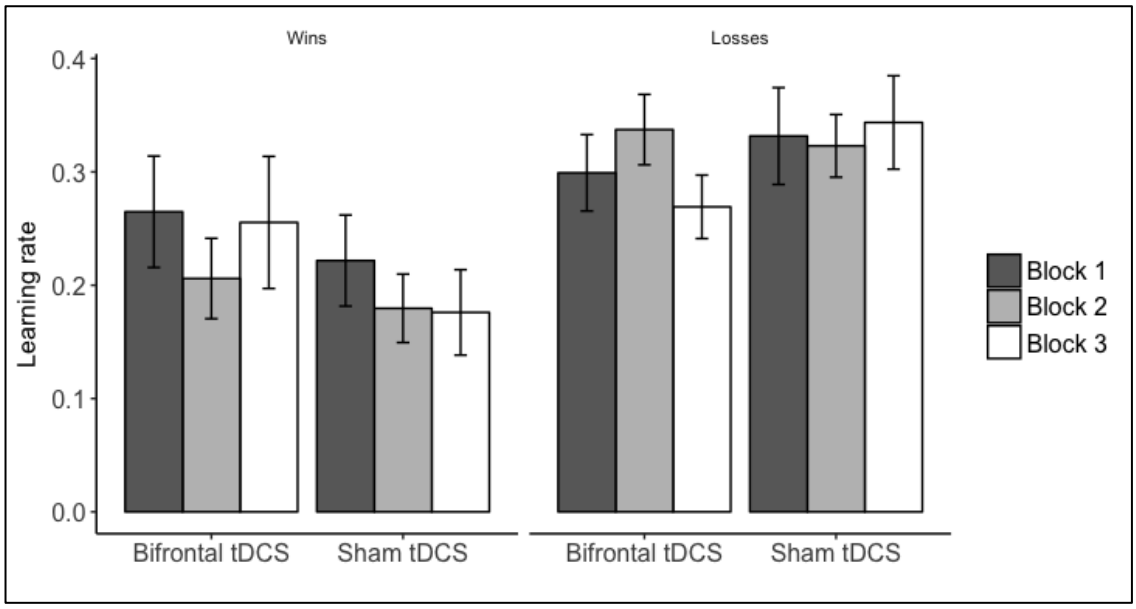

Figure 5.4. Learning rates in the Loss-volatile training blocks by tDCS condition and outcome valence. Error bars represent the standard error of the mean.

5.3.3.2. Choice stochasticity

For the inverse temperature, on the other hand, a significant tDCS condition x Block interaction was observed in the Training blocks (F(1,18) = 3.39, p <0.05, η^2 = 0.036). Follow-up analyses indicated that inverse temperature values increased significantly from Block 1 to Block 3 with bifrontal tDCS (t(19) = 2.97, p <0.01), whereas there were no differences between Block 1 and Block 3 with sham tDCS (t(19) = -0.36, p = 0.725). As shown in Figure 5.5., there appeared to be a linear increase in the inverse temperature with active tDCS, indicating a reduction in stochasticity of choice behaviour across the training blocks. In the Both-volatile blocks, however, there was no significant interaction of tDCS condition and Time (F(1,18) = 0.06, p = 0.809), indicating that stimulation-related changes were not sustained following the training blocks.

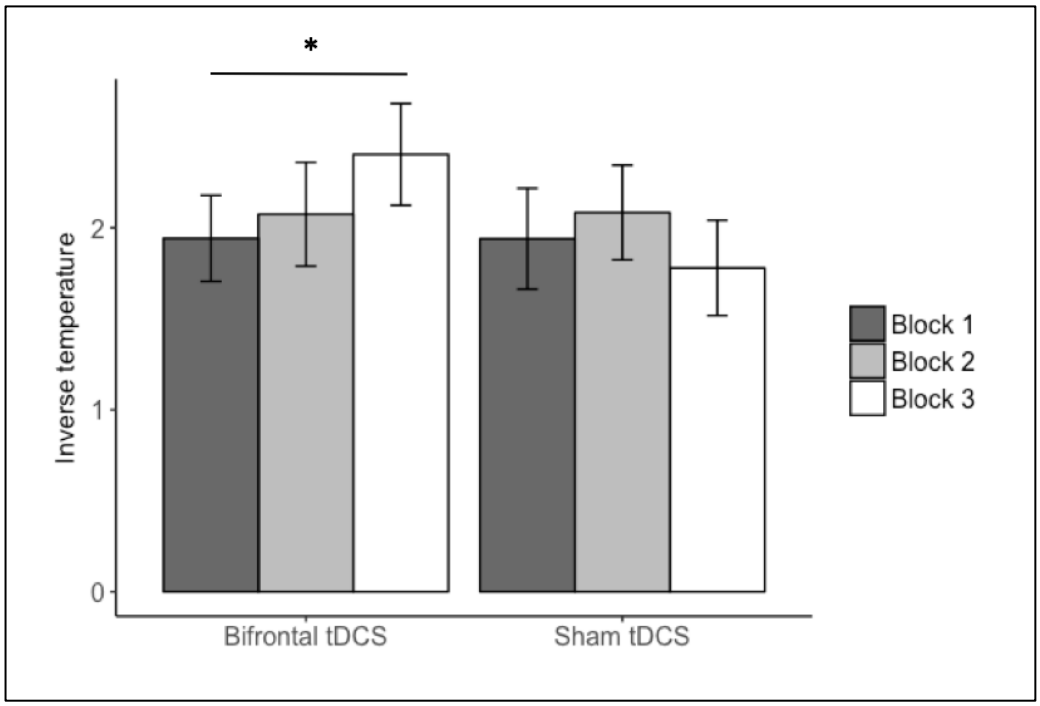

Figure 5.5. Log inverse temperature values in the Loss-volatile training blocks by tDCS condition and Block. Error bars represent the standard error of the mean (*p <0.05).

5.3.3.3. Win-driven choice behaviour and total winnings

For win-driven choice behaviour, there was a significant main effect of Block in the Training blocks (F(2,36) = 4.79, p <0.05, η^2 = 0.020). Paired t-tests indicated that the proportion of win-driven choices was significantly higher in Block 1 than in Block 2 (t(39) = 2.81, p <0.01), while there were no differences between Block 1 and Block 3 (t(39) = 1.39, p = 0.172) or Block 2 and Block 3 (t(39) = -0.95, p = 0.348). There were no further main or interaction effects of Block or tDCS condition (all p >0.05). In the Both-volatile blocks, there was a main effect of tDCS condition (F(1,18) = 4.77, p <0.05, η^2 = 0.059), with the proportion of win-driven choices being greater in the sham tDCS session (M = 0.51, SE = 0.02) than in the active tDCS session (M = 0.44, SE = 0.03). This suggests participants had a greater tendency to base their decisions on the Loss outcomes in the active tDCS session compared to the sham tDCS session. In addition, there was a main effect of Time (F(1,18) = 10.06, p <0.01, η^2 = 0.090). As depicted in Figure 5.6.,

participants tended to choose the win-driven option more frequently before the training blocks (M = 0.52, SE = 0.02) than after the training blocks (M = 0.43, SE = 0.02). This implies that Loss-volatile training might reduce learning from positive outcomes.

The observation of a change in win-driven choices was unexpected given the absence of alterations in learning rates. As the win-driven choice outcome does not distinguish between the influence of different computational parameters, one possible explanation may be that the significant change represents a combination of changes in the learning rate and the inverse temperature. A linear mixed model was run to investigate the relationship between win-driven choice behaviour, learning rates, and the inverse temperature. As shown in Table 5.2., higher learning rates were associated with increases in win-driven choice behaviour, whereas greater loss learning rates were related to decreases in win-driven choice behaviour. In contrast, there appeared to be no significant relationship between the inverse temperature and the proportion of win-driven choices. These findings suggest that the observed changes in win-driven choices with IBLT training are more closely related to the learning rates than to the inverse temperature. Crucially, there was no interaction effect of tDCS condition and Time ($F(1,18) = 1.79$, $p = 0.198$), demonstrating that tDCS did not play a causal role in changing win-driven behaviour over time across the sessions.

Finally, there was no main effect of tDCS condition on the total amount of money won across the Training blocks ($F(1,18) = 0.06$, $p = 0.914$). In addition, there were no main effects of either tDCS condition or Time on money won in the Both-volatile blocks (all $p > 0.05$). Furthermore, there was no interaction of tDCS condition and time on total winnings in the Both-volatile blocks ($F(1,18) = 0.00$, $p = 0.965$).

Table 5.2. Linear mixed model of computational parameter values and win-driven choice behaviour in the Both-volatile blocks of Loss-volatile training.

Predictor	**Coefficient β**	**SE (β)**	***t***	***p***
tDCS condition	0.008	0.013	0.576	0.567
Time	-0.021	0.014	-1.523	0.134
Win learning rate	0.120	0.010	11.674	<0.001
Loss learning rate	-0.076	0.005	-16.82	<0.001
Inverse temperature	-0.001	0.018	0.436	0.664

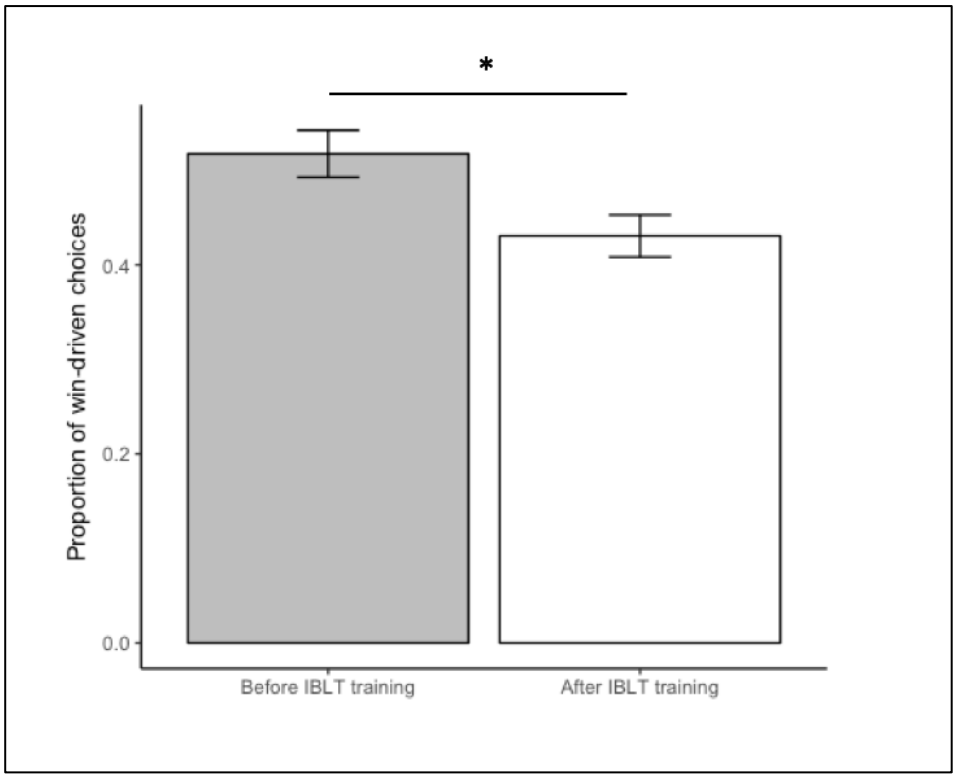

Figure 5.6. Proportion of win-driven choices in the Both-volatile blocks before and after Loss-volatile training. Error bars represent the standard error of the mean (*p <0.01).

5.3.3.4. Facial Expression Recognition Task

On the FERT, significant effects of Emotion were found for task Accuracy (F(5,90) = 10.53, p <0.001, η^2 = 0.248), Reaction time (F(5,90) = 14.08, p <0.001, η^2 = 0.227), and number of Misclassifications (F(5,90) = 7.20, p <0.001, η^2 = 0.211). Similar to the findings reported in the Win-volatile training, participants tended to be more accurate and faster when responding to happy faces then when presented with other emotional expressions.

In addition, participants misclassified happy expressions less frequently compared with the other emotions (see Figure 5.7.). There were no significant main effects of tDCS condition, nor any interactions between tDCS condition and Emotion (all p >0.05).

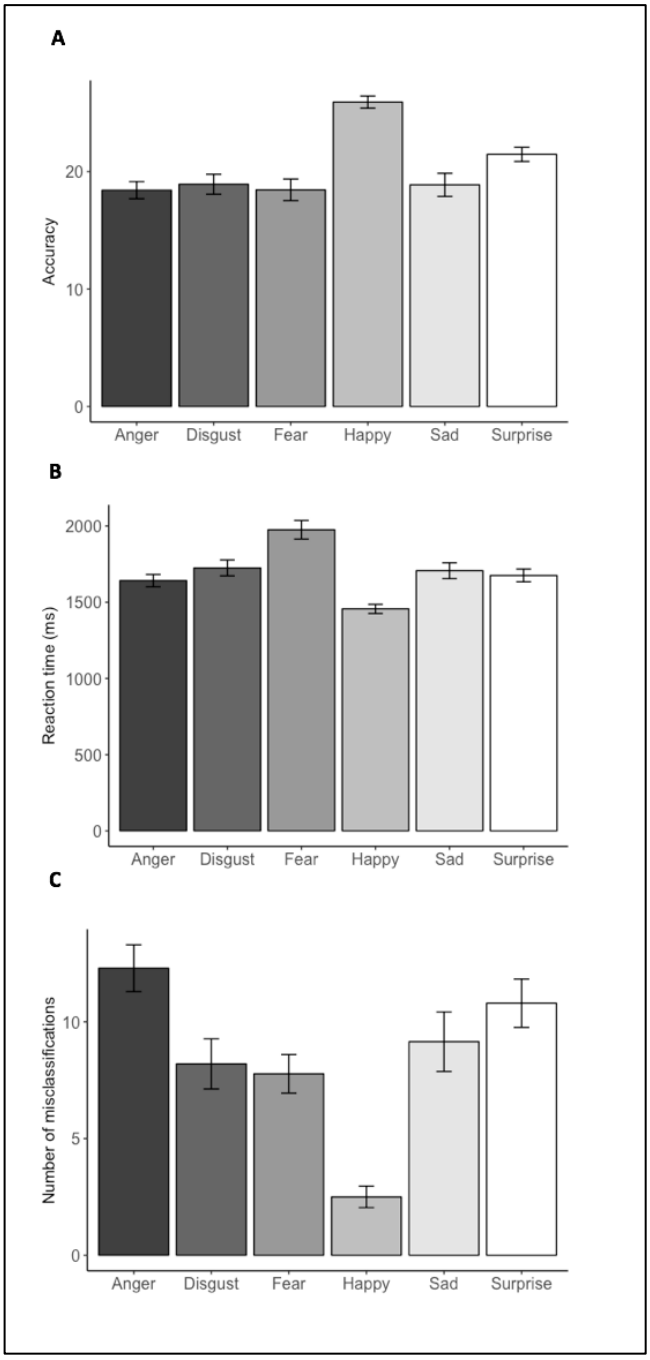

Figure 5.7. Performance on the Facial Emotion Recognition Task for all participants in the Loss-volatile IBLT Training paradigm. **(A)** Number of accurate identifications by emotion. **(B)** Reaction time in milliseconds by emotion. **(C)** Number of misclassifications by emotion. Error bars represent standard error of the mean.

5.3.3.5. Mood and anxiety measures

There were no significant main or interaction effects of Time and tDCS condition on either the PANAS Negative scores or the STAI-State scores (all $p > 0.05$). However, similar to the findings reported for Win-volatile IBLT training, there was a main effect of Time for the PANAS Positive scores ($F(2,38) = 9.23$, $p < 0.001$, $\eta^2 = 0.056$), with participants having higher scores at Time 1 than at Time 2 ($t(39) = 4.24$, $p < 0.001$) and Time 3 ($t(39) = 4.59$, $p < 0.001$), whereas there was no difference between the scores obtained at Time 2 and Time 3 ($t(39) = 0.00$, $p = 0.999$).

5.3.4. Comparison of Win-volatile and Loss-volatile training

5.3.4.1. Learning parameters Training blocks

In order to directly contrast the effects of the two alternate types of training, I conducted additional analyses on the combined data from the Win-volatile and Loss-volatile training paradigms. When investigating changes in learning rates across the Training blocks, a main effect of Outcome valence was observed ($F(1,36) = 36.08$, $p < 0.001$, $\eta^2 = 0.093$), with learning rates being higher for win outcomes (M = 0.31, SE = 0.01) than for loss outcomes (M = 0.22, SE = 0.01). This indicated that participants tended to learn faster from win than loss outcomes independent of the type of training provided, suggestive of a positive learning bias. As shown in Figure 5.8., there was also a main effect of Block ($F(2,72) = 3.45$, $p < 0.05$, $\eta^2 = 0.006$). Specifically, learning rates were higher in Block 1 than in Block 3 ($t(159) = 2.20$, $p < 0.05$), whereas there were no differences between Block 1 and Block 2 ($t(159) = 1.45$, $p = 0.149$) or Block 2 and Block 3 ($t(159) = 1.17$, $p = 0.242$). In addition, there was an interaction of Training type x Outcome valence ($F(1,36) = 147.84$, $p < 0.001$, $\eta^2 = 0.297$). As expected, learning rates were higher for win outcomes than for loss outcomes in the Win-volatile training ($t(119) = 16.02$, $p < 0.001$), whereas the reverse pattern was found in the Loss-volatile training ($t(119) = -6.06$, $p < 0.001$).

Moreover, learning rates for win outcomes were higher in the Win-volatile training than in the Loss-volatile training ($t(238) = 7.72$, $p < 0.001$), while learning rates for loss outcomes were higher in the Loss-volatile training than in the Win-volatile training ($t(237) = 11.21$, $p < 0.001$). These findings indicated that the volatility manipulations had the intended effects on learning rates during the Training blocks. There were no further main or interaction effects of Training type, tDCS condition, Block, or Outcome valence for learning rates or the inverse temperature in the Training blocks (all $p > 0.05$).

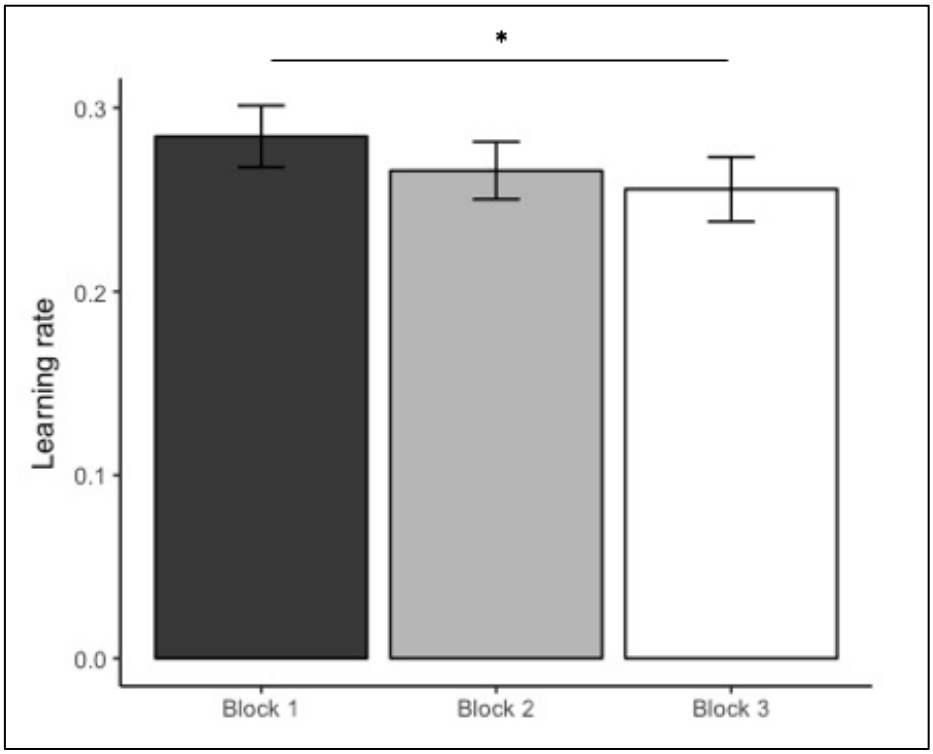

Figure 5.8. Learning rates in the Win-volatile and Loss-volatile training paradigms by training block. Error bars represent the standard error of the mean (*$p < 0.05$).

5.3.4.2. Learning parameters Both-volatile blocks

In the Both-volatile blocks, there was a main effect of Time ($F(1,36) = 16.46$, $p < 0.001$, $\eta^2 = 0.033$), such that learning rates were higher prior to stimulation (M = 0.39, SE = 0.02) than after stimulation (M = 0.29, SE = 0.02). Furthermore, a main effect of Outcome valence was observed ($F(1,36) = 15.21$, $p < 0.001$, $\eta^2 = 0.056$), with participants having higher learning rates for win outcomes (M = 0.38, SE = 0.02) than for loss outcomes (M

= 0.30, SE = 0.02). Finally, there was a Training type x Outcome valence interaction ($F(1,36) = 8.84$, $p < 0.01$, $\eta^2 = 0.033$). It was found that win learning rates were significantly higher than loss learning rates in the Win-volatile training ($t(79) = 5.53$, $p < 0.001$), whereas there was no difference between win and loss learning rates in the Loss-volatile training ($t(79) = 1.04$, $p = 0.30$). In addition, learning rates for loss outcomes were higher in the Loss-volatile training than in the Win-volatile training ($t(157) = 2.91$, $p < 0.01$), while learning rates for win outcomes were similar across the two training types ($t(158) = 1.37$, $p = 0.172$). These findings suggest that participants learned more from loss outcomes in the Loss-volatile than the Win-volatile training study. However, as the Training type x Outcome valence x Time interaction was not significant ($F(1,36) = 1.20$, $p = 0.281$), there is no evidence that this difference was an effect of training rather than a chance difference between the two study samples.

No significant effects were observed for the inverse temperature values in either the Training blocks or the Both-volatile blocks (all $p > 0.05$). To investigate whether the higher baseline inverse temperature values in the Win-volatile training group compared to the Loss-volatile training group affected the outcomes, analyses were re-run using these baseline values as a between-subject covariate. No changes in the outcomes were observed if this factor was included in the analyses.

5.3.4.3. Win-driven choice behaviour

For the proportion of win-driven choices, there was a main effect of Training type in the Training blocks ($F(1,36) = 88.47$, $p < 0.001$, $\eta^2 = 0.596$), with participants making more win-driven choices in the Win-volatile training (M = 0.63, SE = 0.01) than in the Loss-volatile training (M = 0.33, SE = 0.01). In the Both-volatile blocks, a similar main effect of Training type was observed ($F(1,36) = 7.94$, $p < 0.01$, $\eta^2 = 0.089$). Furthermore, there was a main effect of Time ($F(1,36) = 4.19$, $p < 0.05$, $\eta^2 = 0.020$), with the proportion of

win-driven choices being greater before (M = 0.54, SE = 0.02) than after tDCS (M = 0.50, SE = 0.02). In addition, there was a main effect of tDCS condition ($F(1,36) = 6.00$, $p < 0.05$, $\eta^2 = 0.027$), with win-driven choices being more frequent in the sham tDCS condition (M = 0.55, SE = 0.02) than in the active tDCS condition (M = 0.50, SE = 0.02). Although it appeared that the number of win-driven choices decreased more over time with active tDCS compared to sham tDCS (see Figure 5.9.), this effect was not significant ($F(1,36) = 2.19$, $p = 0.148$).

Interestingly, however, there was an interaction of Training type and Time ($F(1,36) = 4.44$, $p < 0.05$, $\eta^2 = 0.021$). Follow-up *t*-tests indicated that the proportion of win-driven choices in the Loss-volatile training was significantly lower after training than before training ($t(39) = 2.96$, $p < 0.01$), while there were no changes over time the Win-volatile training ($t(39) = -0.04$, $p = 0.969$). In addition, there was no significant difference in win-driven behaviour between the two training types before the training commenced ($t(78) = 1.41$, $p = 0.16$), while win-driven choices were made significantly less often in the Loss-volatile group than in the Win-volatile group after training ($t(73) = 3.79$, $p < 0.001$). As shown in Figure 5.10., Loss-volatile training was thus associated with a significant decrease in win-driven behaviour, whereas this type of behaviour was unaltered when engaging in Win-volatile training. There were no further significant effects in either the Training blocks or the Both-volatile blocks (all $p > 0.05$).

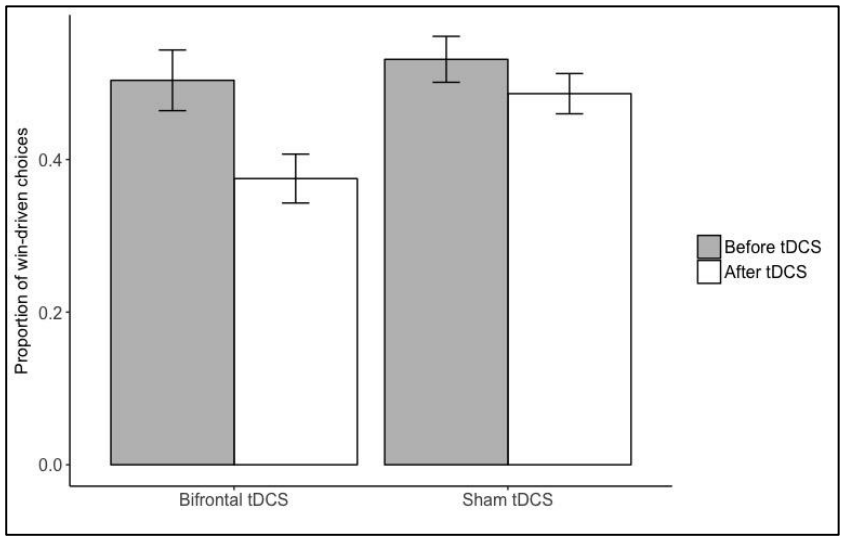

Figure 5.9. Proportion of win-driven choices over time by tDCS condition in the Both-volatile blocks of the Loss-volatile IBLT training paradigm. Error bars represent the standard error of the mean.

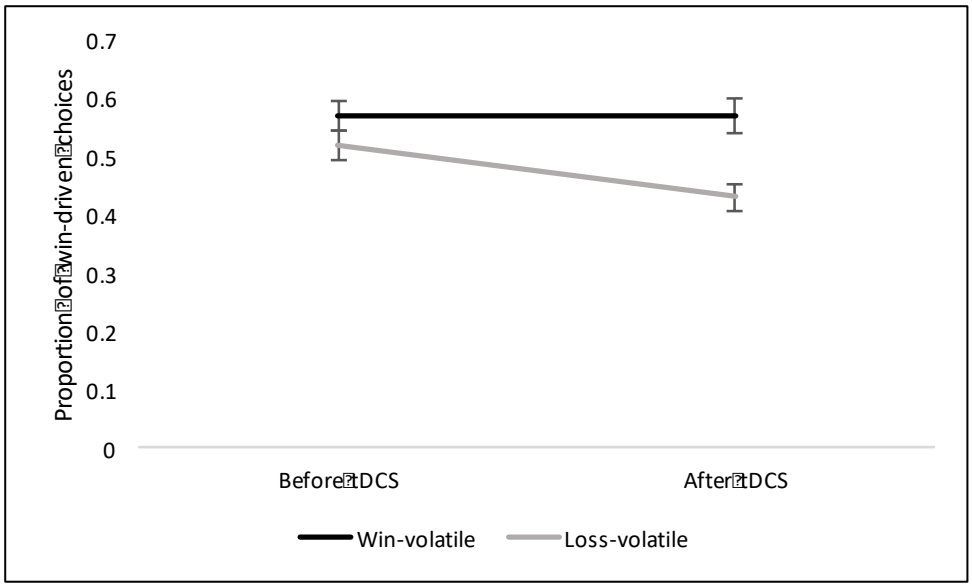

Figure 5.10. Proportion of win-driven choices in the Both-volatile blocks over time by Training type. Error bars represent the standard error of the mean.

5.3.4.4. Facial Expression Recognition Task

Similar to individual analyses of the FERT for the Win-volatile and Loss-volatile training data, main effects of emotion were found for participant's accuracy ($F(5,180) = 21.35$, $p < 0.001$, $\eta^2 = 0.252$), reaction time ($F(5,180) = 31.68$, $p < 0.001$, $\eta^2 = 0.188$), and number

of Misclassifications ($F(5,180) = 19.17$, $p < 0.001$, $\eta^2 = 0.266$). However, there were no interactions of Training type and emotion on any of the outcome variables (all $p > 0.05$), indicating that the two types of IBLT training did not result in differences in emotional processing in healthy adults.

5.3.5. Participant blinding

Across all sessions, participants guessed the correct tDCS condition in 90% of the active tDCS sessions and 33% of the sham tDCS sessions. It should be noted that 25 out of 40 participants chose the active tDCS condition in both sessions, suggesting a lack of knowledge of the true stimulation condition. Interestingly, there was a slight decrease in certainty over time, with participants indicating they were 74.7% confident of their choice in the first session, whereas this value decreased to 70.5% in the second session. Self-reported attention to the IBLT tended to be high, with participants indicating little distraction from the task with either sham tDCS (M = 8.53, SE = 0.21) or active tDCS (M = 8.40, SE = 0.19). This suggests that any differences in task performance between the two tDCS conditions were unlikely to be due to distraction from the task as a consequence of stimulation sensations.

5.4. Discussion

5.4.1. Main findings

The current study investigated the effects of simultaneous cognitive training and bifrontal tDCS on reinforcement learning in healthy adults. In contrast to the hypotheses outlined in the introduction, there were no consistent interaction effects of IBLT training and brain stimulation on learning parameters, attentional biases, or mood in healthy adults. These findings diverge from previous reports indicating that tDCS can amplify the impact of cognitive training (Brunoni et al., 2014a; Clarke et al., 2014; Martin et al., 2013; Segrave

et al., 2014). However, there were various training-specific effects and isolated effects of tDCS on cognitive performance which merit further discussion.

5.4.2. IBLT training effects on reinforcement learning

First of all, it was found that healthy participants were able to adjust their learning behaviour correctly in response to volatility manipulations in both the Win-volatile and Loss-volatile training paradigms. This was demonstrated by the finding that participants in the Win-volatile training group showed high learning rates for win outcomes and low learning rates for loss outcomes across the training blocks, whereas participants in the Loss-volatile training group showed the opposite pattern. Across the two training paradigms, there appeared to be a linear decrease in global learning rates across the blocks. Potentially, this change may have been due to an effect of fatigue over time. Importantly, however, there was no specific progressive change in either win or loss learning rates within the two training paradigms. It can be hypothesised that healthy individuals have a high ability to adapt their behaviour flexibly according to task demands, and do not require further training to improve their performance. This might also explain why neither of the training types was associated with a change over time in the balance of learning rates for rewarding and adverse outcomes in the Both-volatile blocks following training. That is, as the information content for both win and loss outcomes was high in the Both-volatile blocks, participants rapidly adjusted their learning rates to the task demands and applied a high speed of learning for both wins and losses. This could have led to ceiling effects of learning rates in blocks with highly volatile outcomes, reducing the sensitivity to induce or detect potential effects of the stimulation on learning parameters over time. Given the required power estimates for the current sample sizes as reported in Section 4.3.6., the present study was likely underpowered to detect such hypothesised subtle training effects in healthy volunteers. It is possible that greater effects of the training on learning would

be observed in a population in which impairments of adapting learning to volatility changes have been indicated. Alternatively, it may be of use to consider potential adaptations of the IBLT to prevent the occurrence of ceiling effects. For example, it would be possible to introduce an intermediate schedule involving slightly lower volatility of the outcomes, which should result in lower initial learning rates than the present volatile blocks but higher learning rates than for entirely stable outcomes.

In contrast to the null findings from the computational analyses, however, calculations of win-driven behaviour showed that there was a change in the number of win-driven choices made during and following Loss-volatile training. A decrease in win-driven choices was observed from Block 1 to Block 2 during the training itself, although this effect was not maintained in the final training block. After the training had been completed, however, it was found that participants in the Loss-volatile training made fewer win-driven choices compared to pre-training measures on the IBLT. This suggests that there may be some near-transfer effects of Loss-volatile training towards learning from negative outcomes. In the Win-volatile training, on the other hand, no changes were observed in win-driven behaviour in either the Training blocks or the Both-volatile blocks. A potential explanation of the specificity of these effects for Negative- but not Win-volatile training may be that participants had a bias towards processing positive information prior to the training and stimulation, as it was found that participants tended to have higher learning rates for positive than negative outcomes across the two training types. There may therefore have been a smaller margin for increases in win-driven behaviour, whereas there was greater potential for enhancing learning from negative outcomes. This hypothesis of a predisposition towards learning from rewards rather than losses is further supported by the finding that win learning rates in the Loss-volatile training were higher than loss learning rates in the Win-volatile training. Thus, even in the

face of informative negative outcomes, participants attempted to learn from positive outcomes to some degree.

It is unclear why alterations in the win-driven choices were observed whereas closely related learning parameters were unaffected. As the win-driven choice proportion does not differentiate between different internal mechanisms contributing to behaviour, one possibility is that this value reflects a combination of multiple learning components such as the learning rate and the inverse temperature. Further analyses of the relationship between the individual parameters indicated that win-driven choice behaviour was most strongly related to learning rates, whereas there was no clear relation with the inverse temperature. It therefore appears that the observed changes in win-driven choices are more strongly associated with alterations in the speed of learning than alterations in choice stochasticity.

5.4.3. Impact of tDCS on learning parameters

The computational measures used in this study failed to reveal any statistical interactions between tDCS and IBLT training for learning rates in either of the training groups. Contrary to the initial hypotheses, there was no evidence that win learning rates were increased with stimulation of the DLPFC as shown in the study in Chapter 4. One possibility for this lack of change in the Win-volatile group may be that learning rates for positive outcomes quickly reached a ceiling state, such that speed of learning could not be increased further. In addition, the effects of tDCS observed in the previous study indicated that the effects of stimulation on reward learning rates may be limited by a ceiling effect caused by high learning rates when positive outcomes are volatile. The absence of an effect of tDCS in the Loss-volatile training blocks is more puzzling, however, as there were clear indications from the prior study that stimulation tended to increase win learning rates in Loss-volatile blocks (see Chapter 4). One possible explanation may be that the training

towards learning from negative outcomes may have overpowered any effects of tDCS on learning from positive outcomes. Nevertheless, it is surprising that no differences in win learning rates between active and sham tDCS were seen in the first block of training, which may be comparable to the Loss-volatile block carried out in the previous study.

While there were no changes in learning rates across the training sessions with stimulation, findings from the Loss-volatile training group suggested that active tDCS may increase inverse temperature values across the Training blocks compared to sham tDCS. This shows that participants' choices became less random during the training with stimulation of the DLPFC. Bifrontal stimulation may therefore have boosted performance on the IBLT training by reducing randomness of the decisions. However, no changes in inverse temperature were observed after the tDCS and training had ended, implying that the induced alterations were relatively transient. In addition, no effects of tDCS on choice stochasticity were observed in the Win-volatile training. This might mean that the effects of tDCS on the inverse temperature parameter are dependent on the type of training applied. Alternatively, however, the divergent findings across the two training types may be due to differences in baseline values. Prior to training and stimulation, participants in the Loss-volatile group presented with lower inverse temperature values than individuals in the Win-volatile group. Evidence has indicated that individual differences in baseline cognitive ability play a significant role in the effects of both tDCS (Krause & Cohen Kadosh, 2014) and cognitive training (Jaeggi et al., 2014; Katz et al., 2017) on task performance. However, including baseline inverse temperature values in the analyses did not change the observed outcomes, suggesting that initial choice stochasticity was not responsible for the findings reported here. Given that no changes in the inverse temperature were found in any of the other tDCS studies described in this thesis, the increase in inverse temperature in the Loss-volatile training might represent a spurious

finding. Additional research is needed to clarify whether the change observed in Loss-volatile training reflects a chance outcome or a true alteration, and, if the latter, what the cognitive and neural mechanisms underlying the changes in choice stochasticity are.

5.4.4. Affective biases in healthy adults

In both of the training paradigms, it was found that participants tended to be better at recognising happy facial expressions compared to other emotional stimuli. This corresponds with previous findings that healthy adults present with higher accuracy scores and faster responses when shown happy faces compared to facial expressions of other emotions, such as sadness, fear, or disgust (Adams et al., 2016). In this study, responses to visually presented emotions did not change as a result of bifrontal tDCS. Considering the lack of change in learning rates for positive or negative outcomes across the training types, this was not unexpected. However, it also appeared that there were no transfer effects of the different types of training on FERT performance. This was perhaps unsurprising for the Win-volatile training, as increased learning rates from rewards compared to losses appeared to be extinguished immediately after training. In the Loss-volatile training, on the other hand, there were hints that the training had resulted in reduced learning from positive outcomes. It could therefore have been expected that participants receiving Loss-volatile training would be poorer at recognising happiness than those receiving Win-volatile training. As the FERT was not carried out prior to training, however, it is not possible to determine whether there was a change over time in task performance within the Loss-volatile training group.

There are thus no indications of transfer effects of IBLT training to affective biases across the training types. In addition, neither training nor tDCS affected acute mood or anxiety in healthy participants. Once again, this was unsurprising given the absence of changes in learning rates or affective biases. Furthermore, this finding is in agreement with

previous studies suggesting that single sessions of bifrontal tDCS do not alter mood in healthy adults (Morgan et al., 2014; Plazier et al., 2012; Vanderhasselt et al., 2013).

5.4.5. Limitations and future directions

As the present study did not include a sham training condition, it was not possible to examine changes in learning as a function of tDCS in isolation. Potentially, including such a control condition could improve the ability to detect interaction effects of tDCS and IBLT training in healthy adults. Furthermore, the Both-volatile blocks used to measure near-transfer of the training may have been suboptimal given that the main aim of the intervention was to increase rather than lower learning rates for outcomes. As high learning rates for both win and loss outcomes were expected and observed from the Both-volatile blocks from the very first block prior to any training, the potential for increasing learning within these blocks may have been limited. As suggested in section 5.4.2., it may therefore be of use to consider the use of blocks with an intermediate volatility scheme (and thus a slightly lower information content) in future studies to increase sensitivity in measuring sustained increases in learning rates following cognitive training.

Importantly, the null findings reported here were based on data obtained from a population of healthy young adults. The relatively higher learning rates for win outcomes compared to loss outcomes across the IBLT blocks and training types suggest that these individuals were predisposed to focusing on positive information. In addition, as discussed above, the participants were readily able to adjust their learning rates as required by the volatility of the different outcomes. It is plausible that the effects of both training and tDCS would differ in populations in which reinforcement learning is impaired or attention is biased towards negative information. In addition, a single session of concurrent tDCS and cognitive training may have been insufficient to induce changes outlasting the training

itself. In the future, it may therefore be of interest to assess potential alterations in reinforcement learning across multiple training sessions.

5.4.6. Conclusion

In summary, the present study provided no evidence for an enhancing effect of tDCS on training with the IBLT task on learning rates in healthy adults. In addition, there was little evidence that effects of a single session of IBLT training outlasted the training period, although there were some indications that increased learning from negative outcomes might persist in task blocks following the training. However, such changes did not result in any alterations of affective biases in processing facial expressions. Additional research is needed to explore whether this intervention may be more effective in alternative populations.

Chapter 6. Exploring the impact of tDCS over the DLPFC on reward-guided learning in subclinical depression

6.1. Introduction

Major Depressive Disorder (MDD) is a psychiatric illness characterised by persistent low mood and/or loss of pleasure or interest in most activities (American Psychiatric Association, 2013). It has been predicted that the number of people requiring treatment for depression worldwide will rise by 17% to a total of 1.45 million individuals by 2026 (NICE, 2018). MDD has a high impact on individuals' well-being, with depression being the leading cause of disability worldwide (World Health Organisation, 2018) and no more than 3% of all patients experiencing within-normal quality of life prior to treatment (IsHak et al., 2015). In addition, it has been demonstrated that depressive disorders present a significant social and economic burden, with MDD generating a cost of $210.5 billion in 2010 in the United States alone (Greenberg et al., 2015). Currently, the first line of treatment for MDD as recommended by NICE (2018) guidelines includes active monitoring, psychotherapeutic interventions, and/or antidepressant medication, depending on the severity of the disorder. The efficacy of such interventions is generally measured through response and remission rates as indicated by standardised questionnaires and clinical interviews. Achieving remission with the treatments that are currently available is notoriously challenging, with approximately 30% of all patients failing to respond to pharmacological or psychological therapies (Cascalenda, Perry & Looper, 2003; DeRubeis et al., 1999; Rush et al., 2006; Van Hees et al., 2013; Warden et al., 2007). There is therefore an urgent need for the development of novel efficacious interventions for depression.

One potential alternative to classic treatments of depression is the use of non-invasive brain stimulation methods such as transcranial direct current stimulation (tDCS)

and transcranial magnetic stimulation (TMS). Traditionally, such techniques have been deployed to normalise aberrant neural activity associated with general depressive symptoms. The brain region most commonly targeted with tDCS and TMS is the dorsolateral prefrontal cortex (DLPFC), as there is strong evidence to suggest that depressive disorders are related to reduced activation of the left DLPFC, as well as some research indicating a role of hyperactivity of the right DLPFC (Disner et al., 2011; Grimm et al., 2008). Initially, studies of depression tended to focus on the application of TMS over the left DLPFC, which has been shown to result in significant clinical improvements (Kozel & George, 2002). The clinical effectiveness of this protocol has been well-established (George, 2010; George, Taylor, & Short, 2013; Perera et al., 2016), as evidenced by inclusion of prefrontal TMS in NICE guidelines for the treatment of depression (NICE, 2018). Recently, however, there has been a surge of interest in the use of tDCS, which is more portable, less costly, and has been associated with fewer side effects compared with TMS (Brunoni et al., 2012). In tDCS studies on depressive disorders, stimulation is most commonly applied with the anode over the left DLPFC and the cathode over the right DLPFC, with the patient sitting at rest throughout the session. Meta-analyses have indicated that this procedure is associated with significant amelioration of depressive symptoms compared to sham tDCS (Brunoni et al., 2016; Kalu et al., 2012; Meron et al., 2015; Shiozawa et al., 2014), although it has been suggested that effect sizes are small and the clinical utility of tDCS in depression may be limited in its current form (Berlim, Van den Eynde, & Daskalakis, 2015; Kalu et al., 2012). Additional research is thus needed to establish whether the clinical efficacy of tDCS can be enhanced, for example by modification of the stimulation protocol.

As discussed in Chapter 5, an alternate way in which tDCS might be utilised to improve depression may be through modification of specific cognitive processes and

behaviours, such as affective biases, which have been associated with the onset and perpetuation of depressive symptoms. In a study of healthy adults reported in this thesis (Chapter 4), it was attempted to employ bifrontal tDCS in order to modulate reinforcement learning from positive and negative information, a process which may lie at the root of various psychiatric disorders. The main aim of this proof-of-concept study was to enhance activation in the DLPFC and connected regions associated with reinforcement learning, particularly the dorsal anterior cingulate cortex (dACC), and thereby change performance of a probabilistic learning task. It was indicated that bifrontal tDCS compared to sham tDCS may increase the speed of learning from positive but not negative outcomes, particularly when the rewarding outcomes were uninformative. Based on this result, it can be hypothesised that anodal tDCS of the left dorsolateral prefrontal cortex may bias individuals towards learning from and attending to positive information. However, it is presently unclear whether this effect of tDCS on learning in healthy volunteers is generalizable to individuals with depressive symptoms.

It has previously been suggested that the outcome of tDCS may depend on baseline neural activation in brain regions of interest (Krause & Cohen Kadosh, 2014; Kronberg et al., 2017; Woods et al., 2016). Relative to healthy controls, depressed patients tend to present with reduced activation of the left DLPFC (Disner et al., 2011) as well as abnormal activity in the ACC at rest (Dutta, McKie, & Deakin, 2014) and during reward processing (Admon & Pizzagalli, 2015; Holmes & Pizzagalli, 2008; Ubl et al., 2015). Therefore, it is possible that the impact of tDCS over the DLPFC on reinforcement learning performance would differ in people with depression compared to non-depressed individuals. In addition, it has been proposed that baseline cognitive ability may influence the effects of tDCS on task performance (Katz et al., 2017). Previous studies have indicated that learning rates for positive and negative events may be altered in depression compared with healthy

adults, although the exact nature of this change was variable across experimental tasks and studies (Beevers et al., 2013; Chase et al., 2010; Cooper et al., 2014; Dombrovski et al., 2010). While a baseline bias towards learning from rewarding compared to aversive outcomes was observed in healthy volunteers (see Chapter 4), people with depressive symptomatology might present with different initial learning rates for positive and negative events. Alternatively, as proposed in Chapter 1, depression may be characterised by more complex deficits in reward learning. For instance, individuals with high anxiety have been found to adapt their learning rates for negative events in response to changes in the volatility of the environment to a smaller extent than people with low anxiety scores (Browning et al., 2015). It can be conjectured that depressive disorders, which are closely related to symptoms of anxiety, may be similarly associated with impairments in matching the learning rate to the state of the environment. Such baseline differences in reinforcement learning could interact with the impact of bifrontal tDCS to yield different profiles of effect on learning processes across different populations.

In the present pilot study, I therefore explored how tDCS over the DLPFC affected learning parameters in the Information Bias Learning Task (IBLT) in a small sample of individuals with subclinical depression. Specifically, it was investigated whether prefrontal tDCS could increase learning from positive outcomes compared to sham tDCS. This study thus aspired to obtain a first estimate of the generalisability of tDCS effects on learning rates across populations. At a larger scale, the findings of this study may be used to gain further insight into the potential of tDCS as an intervention for affective cognitive impairments associated with depression.

6.2. Methods and Materials

6.2.1. Participant recruitment

The study procedure received CUREC 2 approval from the local university ethics committee (CUREC; R53802/RE001). Participants aged 18 – 45 years were recruited for the study via word of mouth, online advertisements, and posters distributed to university and hospital locations. Exclusion criteria were left-handedness, contraindications to tDCS, use of psychoactive medication, and neurological illness. Time and travel expenses were reimbursed at a rate of £10 per hour, and participants had the opportunity to win a small amount of money through performance of the computerised Information Bias Learning Task (IBLT).

Following provision of online written consent, all potential participants carried out a pre-screening procedure using Qualtrics Software (Qualtrics, Provo, UT), which presented Beck's Depression Inventory II (BDI; Beck, Steer, & Brown, 1996) as well as questions regarding exclusion criteria for tDCS stimulation (see Appendix D). Participants with a minimum score of 14 on the BDI, the cut-off score for mild depression (Beck, Steer, & Brown, 1996), were invited to take part in an in-person screening session conducted at the John Radcliffe Hospital in Oxford. In this session, participants were screened by the researcher for current or previous Axis 1 conditions using the Structured Clinical Interview for DSM-V disorders (SCID Clinical Version; First et al., 2015). Individuals with a personal or familial history of bipolar disorder were excluded from the study, as case studies have indicated that repeated application of tDCS over the DLPFC has the potential to induce manic psychosis or hypomanic episodes (Arul-Anandam, Loo, & Mitchell, 2010; Baccaro et al., 2010; Brunoni et al., 2011b).

In total, fifty-three individuals carried out the online pre-screening procedure as described above. Based on the results of these questionnaires, nineteen potential

participants were then invited to the in-person screening. Thirteen individuals were found to be eligible for the study and were consequently asked to take part in two tDCS sessions. Two participants dropped out of the study following the screening procedure, and one participant decided to discontinue the study after one tDCS session due to persistent tiredness unrelated to the testing session. The final study sample therefore consisted of 10 participants (4 female, mean age = 26.1, SD ± 5.95).

6.2.2. Study design

These ten individuals took part in a within-subject double-blind study, involving one active tDCS session and one sham tDCS session identical to the procedure described in Chapter 4 (Sections 4.2.2 and 4.2.3.). In both of the study sessions, participants were asked to carry out three blocks of the IBLT during the stimulation (Run 1) and after the stimulation had ended (Run 2). In each run of the task, participants completed a block in which win outcomes were volatile and losses were stable (Win-volatile), a block in which loss outcomes were volatile and wins were stable (Loss-volatile), and a block in which both outcomes were volatile (Both-volatile). A detailed description of the Information Bias Learning Task can be found in Section 3.2.3. Analyses were carried out as described in Section 4.2.4.

6.3. Results

6.3.1. Participant characteristics

Depressive symptoms and trait anxiety scores of all participants ranged from 14 to 37 on the BDI (M = 23.5, SE = 3.0) and 39 to 64 on the STAI-Trait (M = 52.4, SE = 3.2), respectively. Eight participants met criteria for Major Depressive Disorder as assessed with the SCID-V. Prior to stimulation, there were no differences in momentary anxiety scores between the bifrontal tDCS session and the sham tDCS session as measured with the STAI-State ($t(9) = 0.17$, $p = 0.869$). However, there was a significant difference

between the two sessions in the PANAS Positive scores ($t(9) = -2.40$, $p < 0.05$), with scores being higher before sham tDCS (M = 24.0, SE = 2.29) than before active tDCS (M = 21.8, SE = 2.02). In addition, PANAS Negative scores tended to be higher before the active tDCS session (M = 15.1, SE = 1.44) than before the sham tDCS session (M = 13.2, SE = 1.18), although this difference failed to reach significance ($t(9) = 1.93$, $p = 0.085$). Overall, these findings indicate that participants presented with a more positive acute mood before the sham tDCS session compared with the active tDCS session. Analyses of the learning parameters were therefore run both with and without PANAS scores as a covariate.

6.3.2. Effects of bifrontal tDCS on learning rates

6.3.2.1. Learning rates Win-volatile and Loss-volatile blocks

The main research question of interest was whether there was a change in learning rates with active bifrontal tDCS in this sample of individuals with subclinical depression. First, the effects of tDCS on learning rates were therefore investigated across the Win-volatile and Loss-volatile blocks carried out during and after stimulation.

There was a main effect of Outcome valence ($F(1,8) = 6.70$, $p < 0.05$, $\eta^2 = 0.029$), which showed that learning rates tended to be higher for win outcomes (M = 0.26, SE = 0.02) than for loss outcomes (M = 0.21, SE = 0.02). In addition, there was a Block x Outcome valence interaction ($F(1,8) = 50.98$, $p < 0.001$, $\eta^2 = 0.190$). It was found that learning rates for wins were higher than learning rates for losses in the Win-volatile blocks ($t(39) = 7.81$, $p < 0.001$). In addition, loss learning rates were higher than win learning rates in the Loss-volatile blocks ($t(39) = 2.70$, $p < 0.05$). Finally, loss learning rates were significantly higher in the Loss-volatile blocks than in the Win-volatile blocks ($t(39) = 6.40$, $p < 0.001$), whereas the opposite pattern was observed for win outcomes ($t(39) = 3.20$, $p < 0.01$). Similar to healthy adults (Chapter 3 and 4), it thus appeared that individuals with symptoms of depression tended to learn faster from volatile than stable outcomes.

Furthermore, there was a tDCS condition x Run interaction ($F(1,8) = 12.18$, $p < 0.01$, $\eta^2 = 0.041$). As shown in Figure 6.1., learning rates were significantly higher in Run 1 with active tDCS than sham tDCS ($t(39) = 3.03$, $p < 0.01$), whereas there was no difference between the two conditions in Run 2 ($t(39) = -0.67$, $p = 0.504$). There was no significant difference in learning rates between Run 1 and Run 2 in either the active tDCS condition ($t(39) = 1.60$, $p = 0.117$) or the sham tDCS condition ($t(39) = -1.71$, $p = 0.096$). It thus appeared that active tDCS was associated with higher global learning rates than sham tDCS during but not after stimulation.

Crucially, the tDCS condition x Outcome valence ($F(1,8) = 0.63$, $p = 0.450$) interaction was not significant across the two Runs. In addition, there were no further significant main effects or interactions if the two Runs were analysed separately (all $p > 0.05$) These results indicate that, contrary to the findings of a modulation of win learning rates with tDCS in Chapter 4, bifrontal stimulation was associated with a general rather than outcome-specific increase in learning rates during stimulation in the present study. Inclusion of baseline win-loss learning rates, PANAS Positive and/or PANAS Negative scores in the analyses did not change the results.

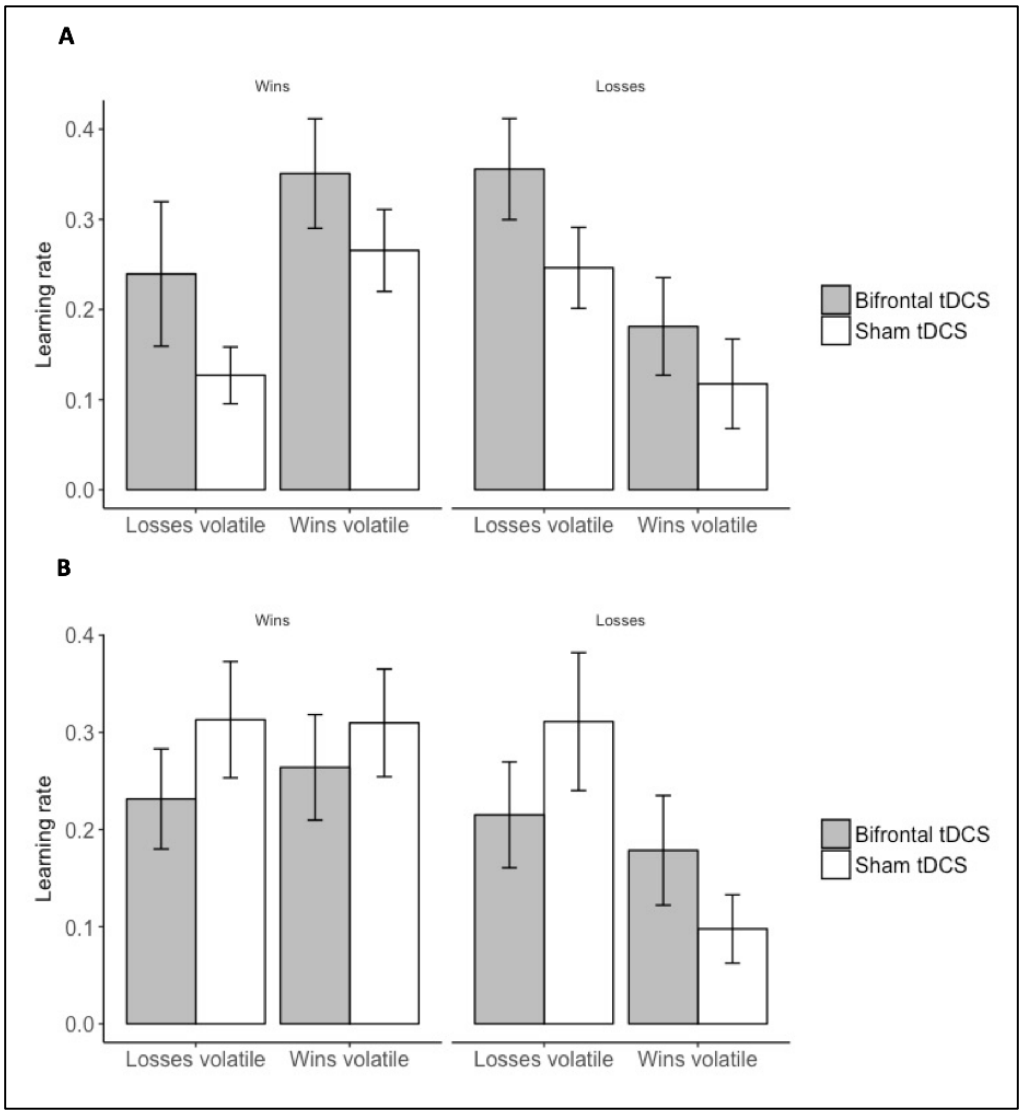

Figure 6.1. Learning rates in the Win- and Loss-volatile block by tDCS condition and Outcome valence. **(A)** Learning rates during tDCS (Run 1) were significantly higher with active than sham stimulation. **(B)** Learning rates after tDCS (Run 2) did not significantly differ as a function of tDCS condition. Error bars represent the standard error of the mean.

6.3.2.2. Learning rates Both-volatile blocks

For the learning rates in the Both-volatile blocks, there were no significant main or interaction effects of tDCS condition, Time, or Outcome valence (all $p > 0.05$). No changes in the results were found if PANAS Positive and/or PANAS Negative scores were included in the analyses.

6.3.3. Effects of bifrontal tDCS on the inverse temperature

To assess potential changes in choice stochasticity with stimulation, participants' inverse temperature values were examined across the tDCS sessions. In the Win- and Loss-volatile blocks, there was a tDCS condition x Run interaction ($F(1,8) = 14.38$, $p < 0.01$, $\eta^2 = 0.067$). Further analyses showed that the inverse temperature tended to be lower with active tDCS than with sham tDCS in Run 1 ($t(39) = -2.92$, $p < 0.01$), whereas the reverse outcome was found in Run 2 ($t(39) = 2.75$, $p < 0.01$) as shown in Figure 6.2. Furthermore, the inverse temperature was significantly higher during stimulation than after stimulation in the sham condition ($t(39) = 3.12$, $p < 0.01$), while the opposite pattern emerged with active tDCS ($t(39) = -2.19$, $p < 0.05$). This indicates that choice stochasticity was greater with active than sham tDCS during stimulation, whereas after stimulation choices were less random in the active than in the sham condition. No further significant main or interaction effects were observed for the inverse temperature parameter in either the Win- and Loss-volatile blocks or the Both-volatile blocks (all $p > 0.05$).

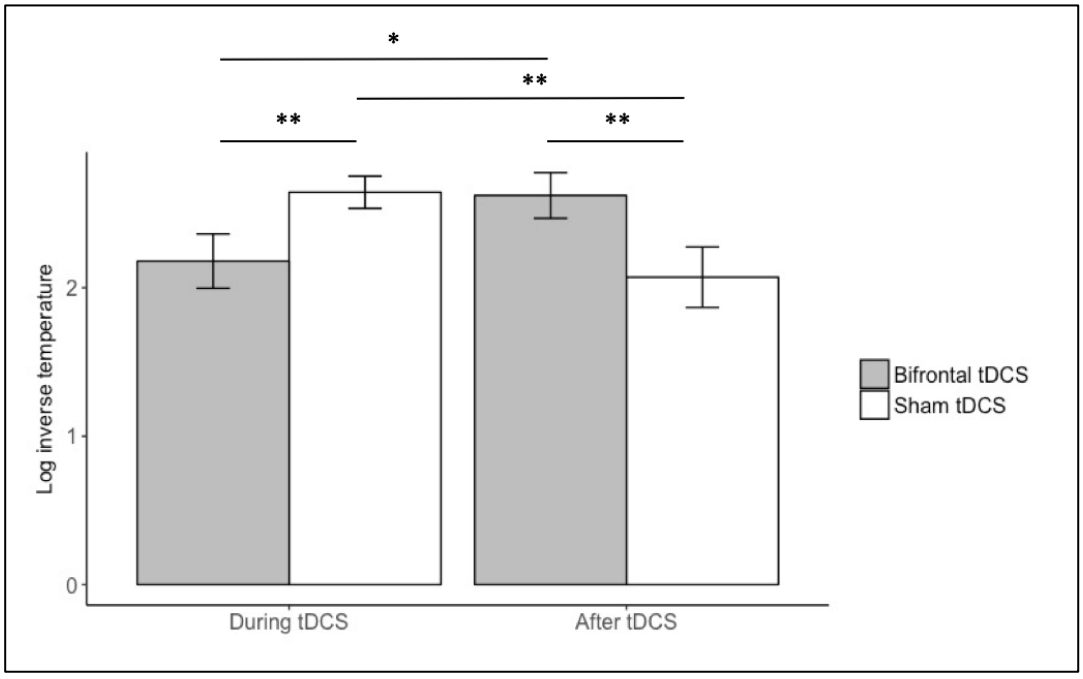

Figure 6.2. Log inverse temperature values on the Win-volatile and Loss-volatile blocks by tDCS Condition and Run. Error bars represent the standard error of the mean (*p <0.05, **p <0.01).

6.3.4. Win-driven choice behaviour and total winnings

In addition to examining the computationally derived parameters, potential changes in win-driven choice behaviour as described in Section 3.2.5. were investigated. In the Win- and Loss-volatile blocks, a main effect of Block was found ($F(1,8) = 292.19$, $p < 0.001$, $\eta^2 = 0.594$), with the proportion of win-driven choices being greater in the Win-volatile blocks (M = 0.67, SE = 0.02) than in the Loss-volatile blocks (M = 0.42, SE = 0.02) similar to findings reported in healthy adults (Chapter 4). Furthermore, there was a Block x Run interaction ($F(1,8) = 26.00$, $p < 0.001$, $\eta^2 = 0.225$), with win-driven choice behaviour being lower in Run 1 (during stimulation) than in Run 2 (after stimulation) in the Loss-volatile block ($t(19) = 5.15$, $p < 0.001$) but not the Win-volatile block ($t(19) = 1.51$, $p = 0.149$). In other words, win-driven choice behaviour increased over time in the Loss-volatile but not the Win-volatile block. As shown in Figure 6.3., there were no significant main effects or interactions of tDCS condition across the Win- or Loss-volatile blocks (all $p > 0.05$).

For the Both-volatile blocks, a main effect of Time was observed ($F(1,8) = 10.73$, $p < 0.05$, $\eta^2 = 0.150$), such that the proportion of win-driven choices was greater at baseline in the first task block before either sham or active tDCS (M = 0.61, SE = 0.04) compared to the last task block (M = 0.45, SE = 0.05). The finding of increased win-driven choice behaviour in Loss-volatile blocks with active tDCS in healthy adults (Chapter 4) was therefore not replicated in the present study. This might be because the proportion of win-driven choices represents a cruder measure of responsiveness to win and loss outcomes than the computational parameters, and may therefore not have been sufficiently sensitive to detect alterations in behaviour with tDCS in the present small study sample. Finally, there was no difference in total winnings between the sham tDCS and active tDCS sessions ($t(9) = -0.33$, $p = 0.752$).

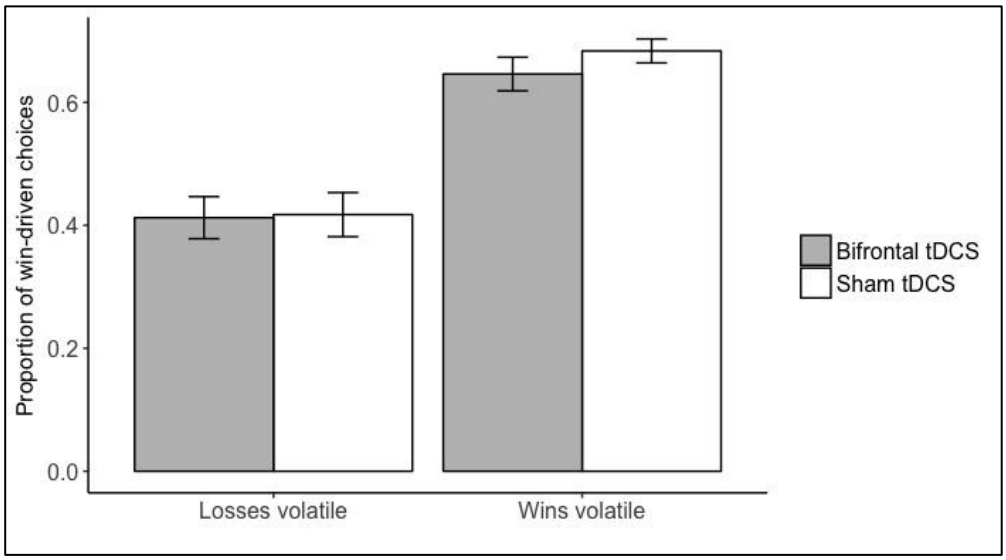

Figure 6.3. Proportion of win-driven choices in the Win-volatile and Loss-volatile blocks by tDCS condition. Error bars represent the standard error of the mean.

6.3.5. Mood and anxiety measures

Considering previous findings of the impact of bifrontal tDCS on depressive symptoms, acute alterations in mood and anxiety associated with the stimulation conditions were also assessed. Baseline questionnaire scores were included as a between-subject covariate to control for the potential influence of differences in baseline scores. There was no significant effect of tDCS condition on change over time in the STAI-State ($F(1,9) = 0.18$, $p = 0.678$) and PANAS Negative scores ($F(1,9) = 0.52$, $p = 0.489$). However, there was a main effect of tDCS condition on change in PANAS Positive scores ($F(1,9) = 6.90$, $p < 0.05$, $\eta^2 = 0.183$). Specifically, there was a greater decrease in positive affect scores over time with sham tDCS (M = -4.0, SE = 0.84) than with active tDCS (M = -1.20, SE = 0.975). As shown in Figure 6.4., this may be due to participants presenting with numerically higher PANAS Positive scores prior to the sham tDCS session than the active tDCS session.

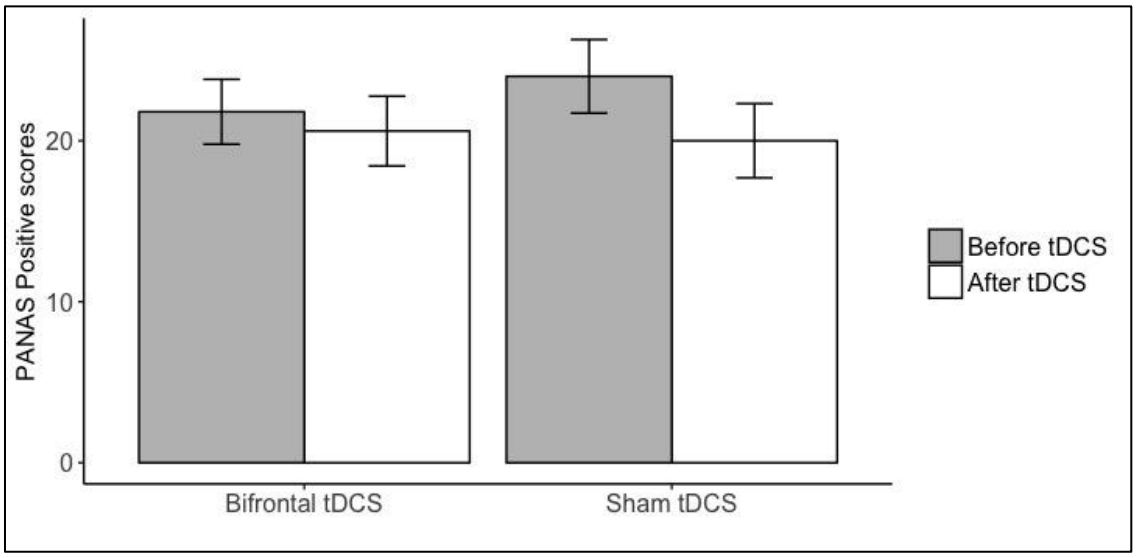

Figure 6.4. PANAS Positive scores before and after stimulation in the sham and active tDCS conditions. Error bars represent the standard error of the mean.

6.3.6. Power calculation

G*Power 3.1 (Faul et al., 2007) calculations were carried out to estimate the required effect size for the present study. It was found that an effect size of 0.996 would be needed to achieve a 0.80 power level for paired comparisons in this sample of 10 participants. This calculation highlights that the current investigation was likely underpowered to reliably detect small to moderate effects of tDCS on learning behaviour.

6.3.7. Participant blinding

Finally, in order to examine whether any changes in behaviour associated with tDCS across the measures could be caused by placebo effects, blinding of the participants across the two testing sessions was assessed. Blinding to the tDCS condition appeared to be successful, with participants identifying the tDCS condition correctly 55% of the time. It was found that 8 out of 10 participants guessed that they had received active stimulation in both sessions, despite being informed that they would receive one sham and one active tDCS session. This suggests that the sham tDCS condition effectively produced stimulation sensations similar to active tDCS. Average certainty of participants regarding

their guess of the tDCS condition was unchanged from the first tDCS session (61%) to the second tDCS session (60%) indicating that one previous experience of the stimulation did not affect participants' subjective ability to distinguish between the two conditions. Moreover, subjective distraction from the task due to the sensations associated with stimulation was similar in the sham tDCS and active tDCS sessions ($t(9) = -1.00$, $p = 0.343$). The effects of tDCS reported here are therefore unlikely to be due to unblinding of the participant or reduced attention caused by stimulation sensations in the active tDCS condition.

6.4. Discussion

6.4.1. Main findings

Based on the preliminary findings of increased reward learning rates with active bifrontal tDCS in healthy adults, as reported in Chapter 4, the present study aimed to assess whether the same stimulation protocol could affect reinforcement learning in a group of individuals with subclinical depression. Overall, it appeared that participants with elevated symptoms of depression were able to adjust their learning rates to the volatility manipulations in the IBLT. In contrast with the results from healthy volunteers, participants with symptoms of depression did not show a specific change in win learning rates with bifrontal tDCS in comparison with sham stimulation. However, it was found that active tDCS was associated with a global increase in learning rates during stimulation compared to the sham condition. The increase in learning rates with active tDCS was not maintained in this sample once stimulation had ceased. There were no indications that tDCS affected win-driven choice behaviour in any of the blocks, which further supports the hypothesis that the balance of learning from win and loss outcomes was unchanged with bifrontal stimulation in the current study sample.

Finally, it was found that bifrontal tDCS altered choice stochasticity across the two

IBLT runs. That is, the inverse temperature was lower with active tDCS than sham tDCS during stimulation, whereas the reverse pattern was found after stimulation. Similar to these findings, it has previously been reported that anodal stimulation of the left DLPFC increased randomness of choice during tDCS in healthy adults (Turi et al., 2015). It can therefore be hypothesised that the left DLPFC may play a role in choice stochasticity, although not all studies have observed changes in the inverse temperature with stimulation of this region (see Hämmerer et al., 2016). Furthermore, it is presently unclear why the inverse effect appeared after tDCS was completed in the current study. Additional research involving larger samples is needed to more clearly determine the impact of bifrontal tDCS on choice stochasticity over time in depressed individuals.

6.4.2. Study limitations

A main limitation of the present study was the small sample size, which reduced the power to identify effects of the stimulation on learning parameters. As discussed in Section 6.3.6., only very high effect sizes would therefore be reliably detected in the current sample. In addition, no upper limit was used for BDI scores during participant recruitment, which resulted in a wide range of scores with the sample. Previously, it has been shown that the effects of bifrontal tDCS on clinical symptoms differ between patients with mild or moderate versus severe depression (Ferrucci et al., 2009). Potentially, effects of stimulation on reward learning mechanisms could similarly differ between individuals with variable scores on indices of depression. While a spread of depression scores would be of use in assessing the influence of severity of depressive symptoms in larger samples, it is possible that this heterogeneity further decreased the power to detect significant effects of tDCS on learning rates in the present small study. Finally, the absence of a healthy control group precluded direct comparisons of the effects of tDCS in depressed and non-depressed individuals.

6.4.3. Conclusion

The present study provided no evidence for specific changes in reward learning rates with bifrontal tDCS in a group of individuals with subclinical depression. However, there was a suggestion that active tDCS increased the overall speed of learning on the IBLT during stimulation, independent of the valence of the outcomes. In addition, it was indicated that choice stochasticity may change over time with active compared to sham bifrontal tDCS of the DLPFC. Recruitment of additional participants is necessary to establish whether these alterations in learning parameters with tDCS in depressed individuals are reliable.

Should anodal stimulation of the left DLPFC indeed be associated with global increases in learning rates, future research might examine whether this effect can be channelled to specifically enhance learning in reward-based tasks. This is of particular interest for the treatment of depression, as it can be hypothesised that changing learning from positive outcomes might influence affective biases associated with mood disorders. In order to optimise potential tDCS interventions, however, it will first be necessary to clarify how reinforcement learning processes are affected in depression. While several previous studies have assessed learning rates for positive and negative outcomes, results in depressed individuals have been mixed. Furthermore, the potential presence of more complex impairments of reinforcement learning, for example learning rate adaptation to stable and volatile environments, has not been examined.

While further research is thus needed to clearly identify reinforcement learning patterns associated with depression, the present findings provide preliminary evidence of effects of bifrontal stimulation on learning mechanisms in individuals with depressive symptoms. In the future, these findings may contribute to the development of tDCS as an interventional technique in depression.

Chapter 7. General discussion

7.1. Summary of results

This thesis aimed to explore the effects of bifrontal transcranial direct current stimulation (tDCS) on reinforcement learning mechanisms in both healthy adults and individuals with symptoms of depression. The experimental studies were designed to answer four main questions: (1) is the performance of healthy adults on a task of reinforcement learning consistent across multiple testing sessions?; (2) can tDCS be applied to alter learning of probabilistic positive or negative stimulus-outcome associations in healthy volunteers?; (3) does tDCS of the dorsolateral prefrontal cortex (DLPFC) augment the effects of cognitive training towards reward or punishment learning?; and (4) can bifrontal tDCS change reinforcement learning processes in a population of individuals with subclinical depression? In the following section, I will discuss the main findings from these individual investigations, consider limitations of the reported studies, and outline potential directions for future research.

7.1.1. Study 1: Performance consistency on the Information Bias Learning Task

In the first experimental study described in Chapter 3 (Study 1), I examined test-retest reliability of and practice effects on the Information Bias Learning Task (IBLT), the main reinforcement learning paradigm utilised in the tDCS studies reported in this thesis. Consistent with previous findings (Pulcu & Browning, 2017), healthy adults were able to generate independent learning rates for positive and negative outcomes on this task. Furthermore, participants adjusted their speed of learning to the volatility of the probabilistic associations across the task blocks. From these adaptations in learning rates it may be inferred that healthy adults were capable of estimating the information content of positive and negative outcomes, and weighted these appropriately in order to optimise their decisions.

In terms of test-retest reliability, it was shown that consistency of task performance over repeated measurements differed across the individual task blocks and outcome types. Moderate reliability indices were found for win outcomes in the Win-volatile blocks, whereas win and loss learning rates in the remaining blocks showed low reliability. These findings suggest that the IBLT may not be suitable for longitudinal studies. However, it was emphasised that this task might be of use in other research designs, including interventions of reinforcement learning and/or the application of IBLT as a training paradigm. In addition to contributing to a better understanding of the psychometric properties of the IBLT, the present results have practical implications for the design and interpretation of future studies. First of all, the intraclass correlation coefficient (ICC) values reported here could be used to inform sample size and power calculations of repeated-measures studies (Fleiss, 1986; Shrout, 1998). In addition, an indication of the expected magnitude of measurement errors on the task was given in the form of standard errors of measurement (SEM) values. These values provide an indication of the minimal differences in scores needed over time to indicate a true change in scores (Weir, 2005) and can be applied to assess whether behaviour has altered more as a result of an intervention than expected by errors of measurement alone. Taken together, the ICC and SEM values may therefore be of use in determining sample sizes for repeated-measures studies and facilitate the interpretation of changes in IBLT learning parameters across sessions.

Of relevance to the within-subjects tDCS research studies presented in this thesis, low test-retest reliability increases the risk of Type II errors and could thus have reduced the power to detect changes in task performance. It is therefore possible that the suboptimal reliability of the Information Bias Learning Task contributed to the null results across these studies. As Study 1 assessed the reliability of task performance in the absence of tDCS, however, it is currently unclear whether the effect of brain stimulation itself on learning

rates is similarly changeable. Of note, test-retest reliability appeared to be particularly low for stable rather than volatile stimuli. This may be due to the high expected uncertainty associated with stable outcomes, which could result in essentially random responses and consequently low correlations between scores across testing sessions. It would therefore be worth exploring whether variations on the "stable" conditions, such as lower expected uncertainty or the introduction of minimal volatility (i.e., an intermediate schedule), would be associated with greater task reliability.

As well as characterising test-retest reliability of the IBLT, this study aimed to identify potential practice effects within and across testing sessions. It was found that healthy participants may present with initial difficulties in utilising negative outcomes to the same extent as positive outcomes in guiding their decisions when the two outcomes are equally informative. This may indicate a predisposition towards processing positive information relative to negative information in healthy adults. However, this tendency appeared to be malleable, as learning rates for negative events increased with repeated exposure to the task blocks. This finding suggests that healthy participants can learn to use negative outcomes to a greater extent with practice of the task. Although this may simply reflect improved task performance, it can be speculated that the present version of the IBLT might encourage learning from negative outcomes within healthy volunteers. Vigilance to possible negative effects of repeated use of this task on information processing may therefore be warranted in future studies. Furthermore, the influence of practice effects on learning rates over multiple testing sessions needs to be kept in mind when interpreting findings from repeated-measures experiments. That is, increases in learning rates for loss stimuli across testing sessions could be due to greater experience in performing the task rather than a result of the intervention. In the tDCS studies reported in the present thesis, this possibility was accounted for by randomising the order of the active

and sham tDCS sessions, thereby eliminating the risk of a main effect of practice driving results of the intervention. Furthermore, the effects of tDCS reported in Chapter 4 mainly concerned win rather than loss learning rates, and are therefore unlikely to be due to practice effects alone. Nevertheless, future studies applying the IBLT may consider the use of a more extensive practice block or session of the task prior to introducing interventions in order to rule out confounding practice effects on changes in reinforcement learning.

7.1.2. Study 2: Prefrontal tDCS increases reward learning rates in healthy adults

In the following study (Study 2), I compared the effects of sham tDCS and active tDCS applied over the dorsolateral prefrontal cortex (DLPFC) on performance of the IBLT in healthy adults. It was shown that tDCS applied during task performance ('online' tDCS) but not tDCS applied at rest ('offline' tDCS) increased learning from reward outcomes. This effect appeared to be particularly strong when the win outcomes were uninformative, which may have been due to a ceiling effect in blocks where win outcomes were informative. The finding that online tDCS was more effective in altering cognitive performance than offline tDCS is consistent with evidence for timing-dependent tDCS effects on cognitive training (Martin et al., 2014) and may be used to inform future non-invasive brain stimulation studies aiming to alter learning processes. In a replication study, however, the interaction of outcome and block for learning rates with tDCS failed to reach significance, which may have been due to a lack of power related to the study's sample size (Minarik et al., 2016).

Despite this failure to replicate in an independent sample, pooled data analyses of the two studies provided supportive evidence that bifrontal tDCS compared to sham stimulation was associated with higher learning rates for both informative and uninformative win outcomes across the two samples. The effects of tDCS may thus be

independent of the volatility of the environment, although the effect is more pronounced when the stimulus-reward associations were kept stable. Active bifrontal tDCS might therefore result in faster learning from positive outcomes, which is a promising finding considering the ultimate objective of this research to enhance information processing of positive events. Alternatively, however, this finding could be interpreted as a failure in estimating how informative positive outcomes are, leading to indiscriminate use of informative and uninformative events. In other words, tDCS might interfere with rather than enhance reinforcement learning processes, resulting in inappropriate use of the available information to inform decisions. As the effects of a single session of tDCS on cognitive and behavioural performance tend to be relatively transient (Antal et al., 2004; Iucalano & Cohen Kadosh, 2013), the present findings do not pose a concern in terms of inducing persistent disruptions of reinforcement learning. However, multiple-session studies should be aware of a possibility of producing adverse effects on cognitive functioning.

In addition, it must be emphasised that the observed increase in learning rates could be caused by several factors. A first interpretation is that stimulation of the DLPFC increased estimations of the volatility of the environment. Alternatively, however, learning rates could change as a result of altered estimations of unexpected uncertainty or noise (see Chapter 1) of the outcomes. In the "stable" conditions, for example, an outcome has a 50% chance of being associated with either of the stimuli and therefore has a high expected uncertainty. When outcomes are highly random, optimal learners should not rely on this information and present with low learning rates. In the "volatile" conditions, on the other hand, an outcome has an 80% chance of being associated with a stimulus and is therefore more predictable. In this case, the expected uncertainty is lower and should therefore lead to higher learning rates. Importantly, it has previously been shown that

learning rates do not vary as a function of expected uncertainty across the IBLT blocks if volatility is kept stable (Pulcu & Browning, 2017). This suggests that expected uncertainty is not responsible for the distinctive pattern of high learning rates for volatile stimuli in the absence of tDCS. However, in the present study it was not possible to establish whether the effects of tDCS on reward learning rates were due to alterations in estimations of expected uncertainty or volatility. It is therefore possible that the increase in learning from uninformative reward outcomes with tDCS reflects a decrease in estimations of the expected uncertainty. Additional research is thus needed to clarify whether tDCS facilitates or interferes with reinforcement learning processes and to disentangle the contribution of volatility and expected uncertainty on these outcomes. Independent of the underlying mechanisms, however, the results demonstrate that tDCS may be applied to specifically increase win learning rates in healthy adults.

7.1.3. Study 3: Training on the IBLT is unaffected by tDCS in healthy adults

Based on the finding that tDCS may increase learning from rewards on the IBLT in healthy volunteers, I next examined whether non-invasive brain stimulation (NIBS) could be applied to boost the effects of a training version of the IBLT on reinforcement learning (Study 3). Two different training types were used to induce a temporary information processing bias. In the Win-volatile paradigm, the training aimed to enhance learning from wins, whereas in the Loss-volatile paradigm training was geared towards increasing learning from losses. The primary questions were whether IBLT training could induce a learning bias in healthy adults, and if bifrontal tDCS altered the effects of the two training paradigms. In addition, I assessed potential transfer of learning to changes in affective biases as a result of concurrent training and tDCS. The rationale was to explore whether a single session of the IBLT combined with prefrontal tDCS might present a novel pathway for changing biased information processing associated with depressive disorders.

In this study, healthy participants received one training session with sham tDCS and one session with active tDCS while performing either the Win-volatile or Loss-volatile IBLT training. The results showed weak indications that Loss-volatile training increased people's tendency to avoid choosing the stimulus which received a loss. This effect was maintained after the training had ended. However, there were no transfer effects of these changes in learning behaviour to measures of either affective biases or acute mood. This showed that the effects of a single session of Loss-volatile IBLT training were limited to the task itself in healthy adults. Moreover, there were no sustained changes in IBLT performance during Win-volatile training with or without neurostimulation, indicating that training towards learning from positive outcomes was ineffective in this group. It is possible that the greater effects observed with Loss-volatile training compared to Win-volatile training were caused by a high baseline propensity of learning from positive events in healthy adults. That is, as these participants tended to have high initial learning rates for win outcomes, the potential for increasing these further with training may have been restricted. In addition, the limited effects of training may reflect the well-developed ability of healthy adults to adjust learning rates to the volatility manipulation of the IBLT. It would therefore be of interest to investigate whether the positive training paradigm applied here has more pronounced effects on learning in people presenting with a negative information processing bias at baseline and/or impairments in adapting learning rates. Alternatively, as suggested in Section 5.4.5., one option might be to utilise training blocks with a lower volatility of win outcomes which could result in lower initial learning rates for positive events and thereby enhance the possibility for change.

While there were thus minor indications that Loss-volatile IBLT training in isolation could increase decision making based on loss outcomes, there was no evidence that bifrontal tDCS could alter performance in either of the training paradigms. It therefore

did not appear to be possible to enhance the effects of training on reinforcement learning by concurrently altering the excitability of task-relevant brain regions within a single training session. One possible explanation for this lack of stimulation effects is that changes might be small when applying a single session of training and tDCS, whereas effects could emerge more strongly over multiple sessions. In previous studies of motor learning, for example, it has been found that active tDCS is associated with greater consolidation of motor skills over days compared with sham tDCS (Reis et al., 2009; Tecchio et al., 2010). It is at present not clear whether such findings extrapolate to cognitive functioning. However, one hypothesis which might be worth exploring further is whether bifrontal tDCS applied during IBLT training could result in enhanced retention of acquired reward or loss learning changes across several training sessions.

Finally, although there were no clear changes in reward or loss learning rates on the IBLT associated with tDCS, the results indicated that a combination of tDCS with training towards learning from negative feedback temporarily reduced randomness of participants' choices. However, it is presently unclear whether this finding reveals a true change in behaviour due to tDCS or reflects a spurious effect, as no further alterations in choice randomness have been observed as a result of bifrontal stimulation in healthy adults in either the training experiments described here or in Study 2. Replication of these effects is therefore needed to validate the robustness of the observed change in inverse temperature values with tDCS.

7.1.4. Study 4: Global changes in learning parameters during tDCS in subclinical depression

Previously, the effects of tDCS have been found to be dependent on baseline cognitive functioning and neural activity (Katz et al., 2017; Krause & Cohen Kadosh, 2014; Li, Uehara, & Hanakawa, 2015). It was therefore unclear whether the impact of brain

stimulation on learning observed in healthy adults (Study 2) would be similar in people with symptoms of depression, which has been associated with hypoactivity of the left DLPFC (Disner et al., 2011; Koenigs & Grafman, 2009) and alterations in learning rates (Beevers et al., 2013; Chase et al., 2010; Cooper et al., 2014; Dombrovski et al., 2010) compared to healthy controls. As described in Chapter 6, I therefore carried out a pilot study examining the impact of tDCS on reinforcement learning in a group of individuals with subclinical depression (Study 4). As this study consisted of a small sample size of 10 individuals, results from this experiment should be interpreted with caution. Of interest, it was found that there was a global increase in learning rates during but not after tDCS compared with sham stimulation. In this sample, there was thus no specific increase in learning from reward outcomes with bifrontal tDCS as observed in healthy adults (Study 2). In addition, it was shown that choice stochasticity was reduced during active stimulation and increased after tDCS, whereas the opposite result was found with sham tDCS. Additional research with greater sample sizes is needed to establish whether the changes in learning parameters with bifrontal stimulation observed in this pilot study are reliable in individuals with elevated symptoms of depression.

7.1.5. Main conclusions

There is an extensive literature investigating the effects of non-invasive brain stimulation over the DLPFC on cognitive processes. Among others, previous research has examined the impact of tDCS on working memory (Martin et al., 2013), impulsive behaviour (Benussi et al., 2017; Shen et al., 2016; Ye et al., 2015), executive functioning (Boggio et al., 2007), and multitasking (Hsu et al., 2015). In addition, there is a large number of studies evaluating the influence of bifrontal tDCS on symptoms of depression (Brunoni et al., 2016; Fregni et al., 2006a; Kalu et al., 2012; Nitsche et al., 2009). However, little research has investigated how tDCS effects might be channelled to alter cognitive

processes directly related to the development and maintenance of depressive disorders. In this thesis, I have reported findings providing the first indications that non-invasive brain stimulation might be able to alter reinforcement learning processes which have previously been associated with depression.

The first proof-of-concept study indicated that reward learning rates may be transiently increased with bifrontal tDCS in healthy adults. It should be noted, however, that this effect was not clearly present in follow-up studies. Firstly, this effect failed to replicate in an independent sample of healthy participants. In addition, there were no indications that active tDCS could increase reward learning rates in a training version of the IBLT in healthy adults. In future research, it would be of interest to explore whether this protocol has more pronounced effects in psychiatric populations or over multiple sessions. Thirdly, there was no evidence for an outcome-specific increase in learning rates in a sample of individuals with subclinical depression. It can therefore be hypothesised that the effects of tDCS on reinforcement learning may depend on population characteristics, although interpretations from this study are curbed by its small sample size.

Further research is thus needed to examine whether it is possible to harness the observed effects of tDCS on learning rates to augment cognitive training or induce specific changes in reward learning and affective biases in depressed individuals, and to assess if such alterations are beneficial to low mood. If this is possible, tDCS might present a novel approach to modulating cognitive processes underlying depressive disorders. Potential future experiments based on the findings reported in this thesis should take into account various methodological aspects and limitations, which are discussed in the following section.

7.2. Methodological considerations

7.2.1. tDCS parameter optimisation

The present studies focused on the application of 2 mA tDCS, with the anode placed over the left DLPFC and the cathode positioned over the right DLPFC. The rationale for the choice of this montage is outlined in Chapter 2. However, it might be worth considering alternative stimulation parameters to alter activity in both the DLPFC and the ACC, and to compare their effectiveness in changing learning rates with effects of the current procedure. For example, studies modelling current distributions with prefrontal tDCS have suggested that placing the cathode over occipital, extracephalic, or cerebellar regions may result in greater activation of the ACC compared to bifrontal tDCS (Bai et al., 2014; Ho et al., 2014). Moreover, increasing the distance between the two electrodes on the head with such montages may increase the relative amount of current entering the brain by minimising "shunting" across the scalp (Datta et al., 2009; Miranda et al., 2006), which can occur in bifrontal stimulation due to the small inter-electrode distance. Research on stimulation of the primary motor cortex, however, has indicated that increasing the electrode distance comes at a cost. That is, when increasing the distance between the electrodes, greater current intensities were needed to induce a similar duration and magnitude of after-effects (Moliadze, Antal, & Paulus, 2010). Although this has not been tested in prefrontal regions, it is possible that similar effects may occur when stimulating the DLPFC. The potential for a trade-off between shunting and current intensity therefore deserves careful consideration and should be examined in more detail. Despite the widespread use of tDCS of the DLPFC in modulating depressive symptoms and cognition, research to date has not directly compared the physiological effects of the different montages. Additional neuroimaging studies are therefore warranted to confirm the influence of different DLPFC montages on activity in the ACC as suggested by

computational models, as well as to determine the efficacy of these alternative parameters in altering reward learning processes.

7.2.2. Computational modelling of the tDCS current

Following recommendations from previous research studies, I modelled the current flow associated with bifrontal stimulation to ascertain whether the applied montage was capable of stimulating the intended brain regions (Bikson, Rahman, & Datta, 2012; De Berker, Bikson, & Bestmann, 2013) using SimNibs 2.0.1. software (Thielscher, Antunes, & Saturnino, 2014). Results were visualised using Gmsh software (Geuzaine & Remacle, 2009) as shown in Figure 7.1.A. It was found that tDCS with the anode over the left DLPFC and the cathode over the right DLPFC, with an electrode size of 25 cm^2 and a current application of 2 mA, induced activations in both the left DLPFC and the right DLPFC. However, relatively higher activation was observed in a region between the two DLPFC, suggesting modulation of activity in the frontopolar cortex (Figure 7.1.B). This is consistent with previous modelling studies, which have indicated that a bilateral DLPFC montage using 7 x 5 electrode pads with a 1 mA current resulted in stimulation of both the left and the right prefrontal cortex (Bai et al., 2013) as well as activation in between the two electrodes (see Figure 7.2.; Rampersad et al., 2014). These findings suggest that the montage applied here was capable of altering activation in the DLPFC, but may also reach alternative prefrontal regions. However, it has been highlighted that direct physiological validation of such models is limited (De Berker, Bikson, & Bestmann, 2013), particularly in the absence of anatomical information from individual participants (Bikson, Rahman, & Datta, 2012). As such, the present results invite further research using neuroimaging methods to provide greater certainty about the region-specific effects of tDCS targeting the DLPFC across individuals.

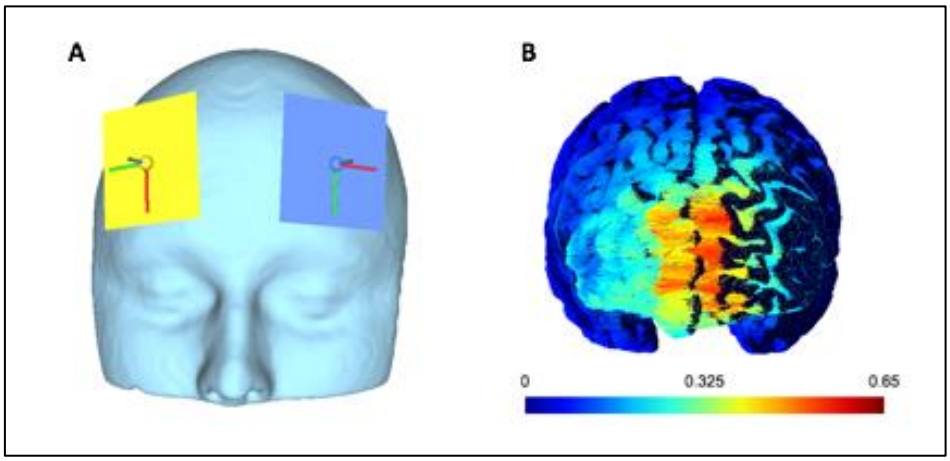

Figure 7.1. (A) Placement of the anodal (blue) and cathodal (yellow) electrodes applied in the present research. **(B)** Simulation of the electric current flow as modelled with SimNibs software, visualised using Gmsh.

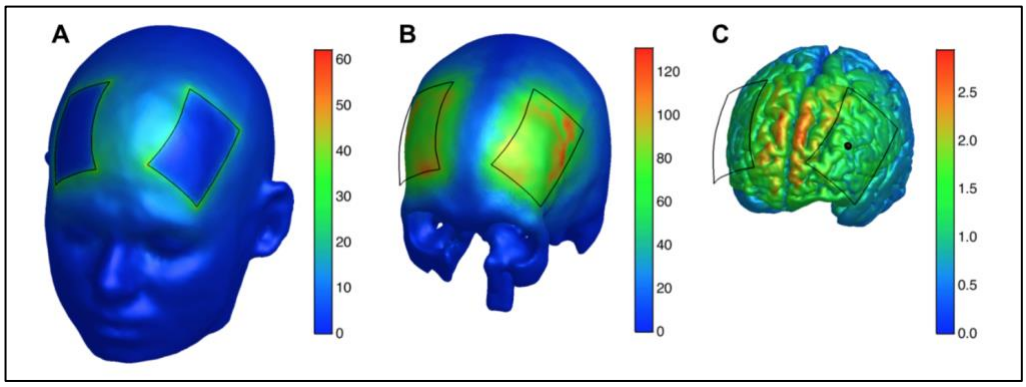

Figure 7.2. Illustration of modelled 1 mA electric field strength on **(A)** the skin surface, **(B)** the skull surface), and **(C)** grey matter when applied over the left DLPFC (anode) and right DLPFC (cathode). Reproduced from Rampersad et al. (2014).

7.2.3. Double-blinding procedure

In addition, one factor which must be taken into account when analysing behavioural changes as a result of non-invasive brain stimulation is the potential for placebo effects. In the present research studies, it was attempted to control for placebo responses by applying a standard double-blinding procedure installed on the tDCS device. However, it has been proposed that participant blinding in tDCS studies may be compromised when utilising a current at 2 mA intensities (O'Connell et al., 2012). More specifically, it has

been found that participants receiving 2 mA tDCS are able to correctly identify the stimulation condition more frequently than would be expected by chance, with confidence in their judgement increasing with repeated administration of the stimulation (O'Connell et al., 2012). In the studies reported here, it appeared that participants were better able to identify active than sham tDCS, with the active tDCS condition being detected correctly in 79% of all sessions and the sham tDCS condition being guessed correctly in 46.4% of all sessions. It is therefore possible that blinding was inadequate, with the potential to induce different expectations of participants across the sessions, which may influence IBLT performance. However, it should be noted that, although participants were explicitly informed that they would receive one session of active tDCS and one session of sham tDCS, 54% of all participants guessed the same condition in both sessions. Moreover, in Studies 3 and 4, certainty rates of judgement were unchanged from Session 1 (72.9%) to Session 2 (71.4%), indicating that repetition of the guessed condition was not due to participants becoming more confident of their knowledge of the condition following the first session.

Furthermore, questioning of the participants following completion of the IBLT (see Appendix F) showed that participants were unaware of the volatility manipulation used in the IBLT and the associated intended effects of the stimulation on their behaviour. In addition, the computationally derived learning parameters of interest are difficult to manipulate intentionally, particularly in the absence of a thorough understanding of how choice behaviour related to learning rates. It is therefore implausible that unblinding of the participants was solely responsible for the observed changes in reward learning rates with tDCS. Moreover, considering the specificity of the effects, it is improbable that distraction from the task with stronger tDCS sensations in the active condition was the cause of the alterations in reinforcement learning. This is further supported by the finding that, when

asked to rate how distracting the stimulation was on a scale of 0 – 10 during task performance, participants' ratings did not differ between the sham condition and the active condition ($t(109) = -0.13$, $p = 0.825$).

7.3. Clinical relevance

While there are thus various limitations which must be considered when evaluating results from the studies reported here, findings from this thesis might be of clinical relevance to the development of novel treatments for depression and other psychiatric disorders. Over the past few decades, a great number of studies have been conducted to examine the effectiveness of tDCS as an intervention for mood disorders, with some success in improving depressive symptoms being observed across studies (e.g. Brunoni et al., 2016; Loo et al., 2012). Compared with traditional treatments of depression, such as antidepressant medication, there are various advantages to the use of tDCS. First of all, tDCS has been rated as a highly acceptable and tolerable intervention by depressed individuals, and may therefore particularly benefit patients who have contraindications or are averse to pharmacological treatments (Meron et al., 2015). In addition, tDCS is relatively easy to apply, portable, and associated with few side effects (Priori, Hallett, & Rothwell, 2009). The recent development of home use devices further broadens the potential clinical applicability of tDCS (Fitz & Reiner, 2013), although it will be essential to develop robust safety standards prior to the use of this technique in private settings.

In previous studies, tDCS has mainly been applied at rest, with the aim of normalising activity in prefrontal brain regions to directly ameliorate general depressive symptoms. In contrast, the present research aimed to alter particular behaviours and cognitions hypothesised to be core features of depressive disorders. Specifically, the goal was to apply non-invasive brain stimulation during the performance of a reinforcement learning task in order to alter reward- and loss-motivated behaviour. A theoretical

advantage of this method is that it targets cognitive mechanisms that have been hypothesised to lie at the root of depressive symptoms, which could potentially result in longer-lasting changes of depression compared with simple modulation of neural activity, which might not be maintained over time. That is, by increasing learning from rewarding outcomes, tDCS might alter information processing and negative affective biases outside of the laboratory, which could then result in sustained improvements of mood regulation. This is of particular interest given the high rates of recurrence of depression in patients, with each depressive episode increasing the likelihood of experiencing further relapses (Kendler et al., 2001). The theoretical basis of this approach is not dissimilar to that of conventional psychotherapeutic treatments of depression such as cognitive behavioural therapy (CBT), which is focused on modifying negative cognitions that are held to be a primary feature of mood disorders (Roiser, Elliott, & Sahakian, 2012). Distinct from CBT, however, the current tactic did not target complex cognitions which are often related to specific social contexts, but rather attempted to modulate basic information processing, which may affect higher-level cognitions as a result. Furthermore, it targeted implicit learning rather than exploring biased cognitions and behaviours explicitly with the participant. A potential benefit of this novel strategy may be that it is somewhat less demanding on neurocognitive resources and less dependent on motivation, which have been suggested to be reduced in patients with depression and can interfere with the effectiveness of psychotherapeutic treatments (Moutoussis et al., 2017; Treadway et al., 2012; Zuroff et al., 2007).

A crucial finding of the present thesis was that bifrontal tDCS could temporarily increase learning from positive outcomes in a reinforcement learning task. This is a promising first step in the development of tDCS as a modulator of cognitive processes associated with depression. Clearly, however, it must be established whether this effect

can be utilised to modify behaviour outside of the task and in ecological settings before its use as a clinical intervention can be considered. In the present research, there was no evidence that tDCS could augment training on the IBLT or change affective biases in healthy adults, suggesting limited transfer of learning changes to other types of stimuli or mood with the procedure in its current form. The training protocol described here thus appeared to be suboptimal for inducing changes in information processing biases in healthy adults. At a conceptual level, however, the present approach of targeting and measuring alterations in computational parameters has the potential to improve upon current non-invasive brain stimulation treatments in depression. First of all, such measures may be more sensitive than assessments of changes in general clinical symptoms. Secondly, analyses of computational parameters can assist in identifying and changing the specific cognitive mechanisms which are responsible for the improvements in depression observed with bifrontal tDCS. Finally, the use of computational models makes it possible to generate and compare more specific hypotheses regarding the effects of tDCS. More extensive use of this method might therefore support clinical programmes in understanding the effects of bifrontal tDCS on the causes of depression, and ultimately contribute to more efficacious treatments with NIBS.

7.4. Future directions

A logical next step within this field would be to assess the effects of tDCS in depressed individuals rather than healthy volunteers. Although a pilot study of changes in learning with tDCS in people with symptoms of depression was carried out here, interpretations of the findings are limited by the small sample size as well as potential differences between subclinical depression and diagnosed Major Depressive Disorder. In addition, the impact of IBLT training combined with tDCS might differ in patients with depression compared to healthy controls, and is therefore worth exploring further. Potentially, training towards

positive outcomes could change learning to a greater extent in depressed individuals, who tend to display a bias towards processing negative information at the cost of processing less positive information (Disner et al., 2011). It should be noted, however, that there has been some uncertainty with respect to how learning rates for positive and negative events are altered in depression with studies using different tasks reporting divergent results (e.g. Beevers et al., 2013; Chase et al., 2010; Cooper et al., 2014; Dombrovski et al., 2010). A clearer understanding of how learning rates are altered in depression, and consequently what the precise aim of the training should be, is therefore of primary importance. Crucially, a research study currently in progress is comparing differences in performance of the IBLT in patients with current depression, individuals with depression, and healthy controls (Browning, 2018). Findings from this study will be critical to the design of optimal IBLT training and tDCS paradigms.

When attempting to induce longer-lasting effects of tDCS on reward learning, for example with the use of repeated stimulation or cognitive training, potential negative outcomes of the neurostimulation procedure should be thoroughly considered. A key finding from the present analyses was that bifrontal tDCS particularly appeared to increase reward learning rates in healthy adults independent of whether the outcomes were informative about the underlying associations. As discussed in Chapter 1, *faster* learning does not necessarily correspond to *better* learning. Rather than interpreting the alterations in learning rate as enhanced learning, it is possible to view the main result as a reduced ability to adapt behaviour to the presented stimulus-outcome associations. As participants receiving active tDCS based their choices in part on an effectively random outcome, which was impossible to predict correctly, they showed suboptimal performance in terms of total winnings (Study 1). It could therefore be suggested that tDCS might impair adaptation of reinforcement learning processes. Furthermore, the consequent effect of repeatedly being

confronted with a negative outcome could potentially confirm negative beliefs present in people with depression. As no alterations in attentional bias or mood were observed in the present studies, it is currently unclear how this increase in reward learning and associated changes in task performance relate to behaviour outside of the IBLT. Future studies should therefore be alert to both positive and negative potential outcomes of bifrontal tDCS, particularly when research involves vulnerable adults.

Finally, one outstanding question is how neural activity related to task performance is altered with transcranial direct current stimulation. Given the ubiquitous use of bifrontal tDCS as an intervention for depression, there has been surprisingly little research to date investigating its effects on cortical excitability. One magnetic resonance spectroscopy study has suggested that active bifrontal tDCS for 30 minutes at 1 mA is related to increased metabolism in the prefrontal cortex and striatum of depressed individuals (Hone-Blanchet, Edden, & Fecteau, 2016). However, this study did not examine potential changes in the anterior cingulate cortex, which would be of particular interest to the present research. Other neuroimaging studies have mainly investigated the effects of prefrontal stimulation with the cathode placed over the right supraorbital region. Resting-state investigations reported increased functional connectivity between the prefrontal cortex and parietal regions (Keeser et al., 2011; Peña-Gómez et al., 2012) as well as the default mode network (Keeser et al., 2011). In addition, an arterial spin labelling study showed that stimulation of the left DLPFC was associated with widespread increases in perfusion across brain regions, including the primary sensory cortex, midcingulate and paracingulate cortex, and left parietal cortex (Stagg et al., 2013). Importantly, it is unclear whether the findings from these studies extrapolate to stimulation with a bifrontal montage, which is commonly used in clinical research on depression. Additional studies are therefore needed to elucidate the physiological effects of bifrontal tDCS. Furthermore, as all of these studies

investigated effects of tDCS applied at rest, the mechanisms underlying effects of online tDCS of the DLPFC have not yet been explored. The introduction of magnetic resonance imaging during stimulation could compare the effects of different montages and verify whether tDCS affects the expected neural regions. In addition, this method could be used to demonstrate which neural changes are most strongly associated with alterations in behaviour. Of particular interest, it has been proposed that model-based fMRI can be applied to correlate individual computational parameters with fMRI signals during task performance (O'Doherty, Hampton, & Kim, 2007). This approach could therefore be applied to tease apart individual changes in neural activity associated with learning rates and inverse temperature values during tDCS as reported in the stimulation studies in this thesis.

7.5. Final remarks

Recent computational studies have highlighted the crucial role of impairments in reward-motivated behaviour in disorders of depression (Huys, Daw, & Dayan, 2015; Moutoussis, Story, & Dolan, 2015; Whitton, Treadway, & Pizzagalli, 2015). In this thesis, it was demonstrated that non-invasive brain stimulation may be able to transiently modulate specific computational components of reinforcement learning in healthy adults. These findings provide preliminary support for the use of tDCS in altering reward learning processes associated with depression. Furthermore, this technique might be of use to other conditions associated with abnormal reinforcement learning processes, including anxiety disorders (Maia & Frank, 2011), schizophrenia (Whitton, Treadway, & Pizzagalli, 2015), obsessive-compulsive disorder (Palminteri et al., 2012), and addiction (Hyman, Malenka, & Nestler, 2006). Importantly, additional research is needed to examine whether transfer of the reported learning changes to affective biases and low mood can be realised, and thereby more firmly establish its potential as a clinical intervention. However, by

indicating the potential of modulating hypothesised causal factors of mood disorders with tDCS, the current results provide an exciting early step in the development of novel treatments for depression.

References

Adams, R. A., Huys, Q. J. M., & Roiser, J. P. (2016). Computational Psychiatry: Towards a mathematically informed understanding of mental illness. *Journal of Neurology, Neurosurgery and Psychiatry*, *87*(1), 53 – 63.

Adams, T., Pounder, Z., Preston, S., Hanson, A., Gallagher, P., Harmer, C. J., & McAllister-Williams, R. H. (2016). Test–retest reliability and task order effects of emotional cognitive tests in healthy subjects. *Cognition and Emotion*, *30*(7), 1247 – 1259.

Admon, R., Nickerson, L. D., Dillon, D. G., Holmes, A. J., Bogdan, R., Kumar, P., ... & Pizzagalli, D. A. (2015). Dissociable cortico-striatal connectivity abnormalities in major depression in response to monetary gains and penalties. *Psychological Medicine*, *45*(1), 121 – 131.

Admon, R., & Pizzagalli, D. A. (2015). Dysfunctional reward processing in depression. *Current Opinion in Psychology*, *4*, 114 – 118.

Aiba, A., Chen, C., Herrup, K., Rosenmund, C., Stevens, C. F., & Tonegawa, S. (1994). Reduced hippocampal long-term potentiation and context-specific deficit in associative learning in mGluR1 mutant mice. *Cell*, *79*(2), 365 – 375.

Akitsuki, Y., Sugiura, M., Watanabe, J., Yamashita, K., Sassa, Y., Awata, S., ... & Kawashima, R. (2003). Context-dependent cortical activation in response to financial reward and penalty: An event-related fMRI study. *Neuroimage*, *19*(4), 1674 – 1685.

Alloy, L. B., Olino, T., Freed, R. D., & Nusslock, R. (2016). Role of reward sensitivity and processing in major depressive and bipolar spectrum disorders. *Behavior Therapy*, *47*(5), 600 – 621.

American Psychiatric Association (2013). *Diagnostic and Statistical Manual of Mental Disorders (5th ed.).* Arlington, VA: American Psychiatric Publishing.

Amiez, C., Joseph, J. P., & Procyk, E. (2006). Reward encoding in the monkey anterior cingulate cortex. *Cerebral Cortex, 16*(7), 1040 – 1055.

Anderson, I. M., Shippen, C., Juhasz, G., Chase, D., Thomas, E., Downey, D., ... & Deakin, J. W. (2011). State-dependent alteration in face emotion recognition in depression. *The British Journal of Psychiatry*, *198*(4), 302 – 308.

Antal, A., Alekseichuk, I., Bikson, M., Brockmöller, J., Brunoni, A. R., Chen, R., ... & Fregni, F. (2017). Low intensity transcranial electric stimulation: Safety, ethical, legal regulatory and application guidelines. *Clinical Neurophysiology*, *128*(9), 1774 – 1809.

Antal, A., Kincses, T. Z., Nitsche, M. A., Bartfai, O., & Paulus, W. (2004). Excitability changes induced in the human primary visual cortex by transcranial direct current stimulation: Direct electrophysiological evidence. *Investigative Ophthalmology & Visual Science*, *45*(2), 702 – 707.

Antonenko, D., Schubert, F., Bohm, F., Ittermann, B., Aydin, S., Hayek, D., ... & Flöel, A. (2017). tDCS-induced modulation of GABA levels and resting-state functional connectivity in older adults. *Journal of Neuroscience*, *37*(15), 4065 – 4073.

Arul-Anandam, A. P., Loo, C., & Mitchell, P. (2010). Induction of hypomanic episode with transcranial direct current stimulation. *The Journal of ECT*, *26*(1), 68 – 69.

Aston-Jones, G., & Cohen, J. D. (2005). An integrative theory of locus coeruleus-norepinephrine function: Adaptive gain and optimal performance. *Annual Review of Neuroscience*, *28*(1), 403 – 450.

Baccaro, A., Brunoni, A. R., Bensenor, I. M., & Fregni, F. (2010). Hypomanic episode in unipolar depression during transcranial direct current stimulation. *Acta Neuropsychiatrica*, *22*(6), 316 – 318.

Bachtiar, V., Near, J., Johansen-Berg, H., & Stagg, C. J. (2015). Modulation of GABA and resting state functional connectivity by transcranial direct current stimulation. *Elife*, *4*, e08789.

Baert, S., De Raedt, R., Schacht, R., & Koster, E. H. (2010). Attentional bias training in depression: Therapeutic effects depend on depression severity. *Journal of Behavior Therapy and Experimental Psychiatry*, *41*(3), 265 – 274.

Bai, S., Dokos, S., Ho, K. A., & Loo, C. (2014). A computational modelling study of transcranial direct current stimulation montages used in depression. *Neuroimage*, *87*, 332 – 344.

Batsikadze, G., Moliadze, V., Paulus, W., Kuo, M. F., & Nitsche, M. A. (2013). Partially non-linear stimulation intensity-dependent effects of direct current stimulation on motor cortex excitability in humans. *Journal of Physiology*, *591*(7), 1987 – 2000.

Bechara A. (2007). *Iowa Gambling Task (IGT) Professional Manual.* Lutz, FL: Psychological Assessment Resources.

Beck, A. T. (1967). *Depression: Clinical, experimental, and theoretical aspects*. New York: Harper and Row.

Beck, A. T. (2008). The evolution of the cognitive model of depression and its neurobiological correlates. *American Journal of Psychiatry*, *165*(8), 969 – 977.

Beck, A. T., Steer, R. A., & Brown, G. K. (1996). Beck depression inventory-II. *San Antonio*, *78*(2), 490 – 498.

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

Beevers, C. G., Clasen, P. C., Enock, P. M., & Schnyer, D. M. (2015). Attention bias modification for major depressive disorder: Effects on attention bias, resting state connectivity, and symptom change. *Journal of Abnormal Psychology*, *124*(3), 463 – 475.

Beevers, C. G., Worthy, D. A., Gorlick, M. A., Nix, B., Chotibut, T., & Maddox, W. T. (2013). Influence of depression symptoms on history-independent reward and punishment processing. *Psychiatry Research*, *207*(1), 53 – 60.

Behrens, T. E., Hunt, L. T., Woolrich, M. W., & Rushworth, M. F. (2008). Associative learning of social value. *Nature*, *456*(7219), 245 – 249.

Behrens, T. E. J., Woolrich, M. W., Walton, M. E., & Rushworth, M. F. S. (2007). Learning the value of information in an uncertain world. *Nature Neuroscience*, *10*(9), 1214 – 1221.

Bennett, M. R. (2000). The concept of long term potentiation of transmission at synapses. *Progress in Neurobiology*, *60*(2), 109 – 137.

Benussi, A., Alberici, A., Cantoni, V., Manenti, R., Brambilla, M., Dell'Era, V., ... & Cotelli, M. (2017). Modulating risky decision-making in Parkinson's disease by transcranial direct current stimulation. *European Journal of Neurology*, *24*(5), 751 – 754.

Berlim, M. T., Van den Eynde, F., & Daskalakis, Z. J. (2013). Clinical utility of transcranial direct current stimulation (tDCS) for treating major depression: A systematic review and meta-analysis of randomized, double-blind and sham-controlled trials. *Journal of Psychiatric Research*, *47*(1), 1 – 7.

Berlim, M. T., Van den Eynde, F., Tovar-Perdomo, S., & Daskalakis, Z. J. (2014). Response, remission and drop-out rates following high-frequency repetitive transcranial magnetic stimulation (rTMS) for treating major depression: A systematic review and meta-analysis of randomized, double-blind and sham-controlled trials. *Psychological Medicine, 44*(2), 225 – 239.

Bhagwagar, Z., Cowen, P. J., Goodwin, G. M., & Harmer, C. J. (2004). Normalization of enhanced fear recognition by acute SSRI treatment in subjects with a previous history of depression. *American Journal of Psychiatry*, *161*(1), 166 – 168.

Bikson, M., Grossman, P., Thomas, C., Zannou, A. L., Jiang, J., Adnan, T., ... & Brunoni, A. R. (2016). Safety of transcranial direct current stimulation: Evidence based update 2016. *Brain Stimulation: Basic, Translational, and Clinical Research in Neuromodulation*, *9*(5), 641 – 661.

Bikson, M., Inoue, M., Akiyama, H., Deans, J. K., Fox, J. E., Miyakawa, H., & Jefferys, J. G. (2004). Effects of uniform extracellular DC electric fields on excitability in rat hippocampal slices in vitro. *The Journal of Physiology*, *557*(1), 175 – 190.

Bikson, M., Rahman, A., & Datta, A. (2012). Computational models of transcranial direct current stimulation. *Clinical EEG and Neuroscience*, *43*(3), 176 – 183.

Bikson, M., Rahman, M., Datta, A., Fregni, F., & Merabet, L. (2013). High-resolution modeling assisted design of customised and individualised transcranial direct current stimulation protocols. *Neuromodulation: Journal of the International Neuromodulation Society*, *15*(4), 306 – 315.

Bindman, L. J., Lippold, O. C. J. & Redfearn, J. W. T. (1962). Long-lasting changes in the level of the electrical activity of the cerebral cortex produced by polarizing currents. *Nature, 196*, 584 – 585.

Blanco, N. J., Otto, A. R., Maddox, W. T., Beevers, C. G., & Love, B. C. (2013). The influence of depression symptoms on exploratory decision-making. *Cognition*, *129*(3), 563 – 568.

Blumberger, D.M., Tran, L.C., Fitzgerald, P.B., Hoy, K.E., & Daskalakis, Z.J. (2012). A randomized double-blind sham-controlled study of transcranial direct current stimulation for treatment-resistant major depression. *Frontiers in Psychiatry, 3*, 74.

Boggio, P. S., Bermpohl, F., Vergara, A. O., Muniz, A. L., Nahas, F. H., Leme, P. B., ... & Fregni, F. (2007). Go-no-go task performance improvement after anodal transcranial DC stimulation of the left dorsolateral prefrontal cortex in major depression. *Journal of Affective Disorders*, *101*(1), 91 – 98.

Boggio, P. S., Campanhã, C., Valasek, C. A., Fecteau, S., Pascual-Leone, A., & Fregni, F. (2010). Modulation of decision-making in a gambling task in older adults with transcranial direct current stimulation. *European Journal of Neuroscience*, *31*(3), 593 – 597.

Boggio, P. S., Ferrucci, R., Rigonatti, S. P., Covre, P., Nitsche, M., Pascual-Leone, A., & Fregni, F. (2006). Effects of transcranial direct current stimulation on working memory in patients with Parkinson's disease. *Journal of the Neurological Sciences*, *249*(1), 31 – 38.

Boggio, P. S., Rigonatti, S. P., Ribeiro, R. B., Myczkowski, M. L., Nitsche, M. A., Pascual-Leone, A., & Fregni, F. (2008). A randomized, double-blind clinical trial on the efficacy of cortical direct current stimulation for the treatment of major depression. *International Journal of Neuropsychopharmacology*, *11*(2), 249 – 254.

Brevet-Aeby, C., Brunelin, J., Iceta, S., Padovan, C., & Poulet, E. (2016). Prefrontal cortex and impulsivity: Interest of noninvasive brain stimulation. *Neuroscience and Biobehavioral Reviews*, *71*, 112 – 134.

Browning, M. (2018, June 11). *Dynamic Learning in Depression.* Retrieved from https://clinicaltrials.gov/ct2/show/NCT02913898.

Browning, M., Behrens, T. E., Jocham, G., O'Reilly, J. X., & Bishop, S. J. (2015). Anxious individuals have difficulty learning the causal statistics of aversive environments. *Nature Neuroscience*, *18*(4), 590 – 596.

Browning, M., Holmes, E. A., Charles, M., Cowen, P. J., & Harmer, C. J. (2012). Using attentional bias modification as a cognitive vaccine against depression. *Biological Psychiatry*, *72*(7), 572 – 579.

Browning, M., Holmes, E. A., & Harmer, C. J. (2010). The modification of attentional bias to emotional information: A review of the techniques, mechanisms, and relevance to emotional disorders. *Cognitive, Affective and Behavioral Neuroscience*, *10*(1), 8 – 20.

Browning, M., Reid, C., Cowen, P. J., Goodwin, G. M., & Harmer, C. J. (2007). A single dose of citalopram increases fear recognition in healthy subjects. *Journal of Psychopharmacology*, *21*(7), 684 – 690.

Brunoni, A. R., Boggio, P. S., De Raedt, R., Bensenor, I. M., Lotufo, P. A., Namur, V., ... & Vanderhasselt, M. A. (2014). Cognitive control therapy and transcranial direct current stimulation for depression: A randomized, double-blinded, controlled trial. *Journal of Affective Disorders*, *162*, 43 – 49.

Brunoni, A. R., Ferrucci, R., Bortolomasi, M., Scelzo, E., Boggio, P. S., Fregni, F., ... & Priori, A. (2013a). Interactions between transcranial direct current stimulation

(tDCS) and pharmacological interventions in the major depressive episode: Findings from a naturalistic study. *European Psychiatry*, *28*(6), 356 – 361.

Brunoni, A. R., Ferrucci, R., Bortolomasi, M., Vergari, M., Tadini, L., Boggio, P. S., ... & Priori, A. (2011a). Transcranial direct current stimulation (tDCS) in unipolar vs. bipolar depressive disorder. *Progress in Neuro-Psychopharmacology and Biological Psychiatry*, *35*(1), 96 – 101.

Brunoni, A. R., Moffa, A. H., Fregni, F., Palm, U., Padberg, F., Blumberger, D. M., ... & Loo, C. K. (2016). Transcranial direct current stimulation for acute major depressive episodes: Meta-analysis of individual patient data. *The British Journal of Psychiatry, 208*, 522 – 531.

Brunoni, A. R., Nitsche, M. A., Bolognini, N., Bikson, M., Wagner, T., Merabet, L., … Bolognini, N. (2012). Clinical research with transcranial direct current stimulation (tDCS): Challenges and future directions. *Brain Stimulation*, *5*(3), 175 – 195.

Brunoni, A. R., Valiengo, L., Baccaro, A., Zanao, T. A., de Oliveira, J. F., Goulart, A., ... & Fregni, F. (2013b). The sertraline vs electrical current therapy for treating depression clinical study: results from a factorial, randomized, controlled trial. *JAMA Psychiatry*, *70*(4), 383 – 391.

Brunoni, A. R., Valiengo, L., Zanao, T., de Oliveira, J. F., Bensenor, I. M., & Fregni, F. (2011). Manic psychosis after sertraline and transcranial direct-current stimulation. *The Journal of Neuropsychiatry and Clinical Neurosciences*, *23*(3), E4 – E5.

Brunoni, A. R., & Vanderhasselt, M. A. (2014). Working memory improvement with non-invasive brain stimulation of the dorsolateral prefrontal cortex: A systematic review and meta-analysis. *Brain and Cognition*, *86*, 1 – 9.

Brunoni, A. R., Zanao, T. A., Ferrucci, R., Priori, A., Valiengo, L., de Oliveira, J. F., … Fregni, F. (2013c). Bifrontal tDCS prevents implicit learning acquisition in antidepressant-free patients with major depressive disorder. *Progress in Neuro-Psychopharmacology and Biological Psychiatry*, *43*, 146 – 150.

Casacalenda, N., Perry, J. C., & Looper, K. (2002). Remission in major depressive disorder: A comparison of pharmacotherapy, psychotherapy, and control conditions. *American Journal of Psychiatry*, *159*(8), 1354 – 1360.

Casula, E. P., Testa, G., Bisiacchi, P. S., Montagnese, S., Caregaro, L., Amodio, P., & Schiff, S. (2017). Transcranial direct current stimulation (tDCS) of the anterior prefrontal cortex (aPFC) modulates reinforcement learning and decision-making under uncertainty: A double-blind crossover study. *Journal of Cognitive Enhancement*, *1*(3), 318–326.

Cavanagh, J. F., Bismark, A. J., Frank, M. J., & Allen, J. J. B. (2011). Larger error signals in major depression are associated with better avoidance learning. *Frontiers in Psychology*, *2*, 1 – 6.

Cella, M., Dymond, S., & Cooper, A. (2010). Impaired flexible decision-making in major depressive disorder. *Journal of Affective Disorders, 124*(1–2), 207 – 210.

Chase, H. W., Frank, M. J., Michael, A., Bullmore, E. T., Sahakian, B. J., & Robbins, T. W. (2010). Approach and avoidance learning in patients with major depression and healthy controls: Relation to anhedonia. *Psychological Medicine*, *40*(3), 433 – 440.

Chatrian, G.E., Lettich, E., & Nelson, P.L. (1985). Ten percent electrode system for topographic studies of spontaneous and evoked EEG activity. *The American Journal of EEG Technology, 25*, 83 – 92.

Chen, C., Takahashi, T., Nakagawa, S., Inoue, T., & Kusumi, I. (2015). Reinforcement learning in depression: A review of computational research. *Neuroscience & Biobehavioral Reviews*, *55*, 247 – 267.

Cheng, G. L. F., & Lee, T. M. C. (2016). Altering risky decision-making: Influence of impulsivity on the neuromodulation of prefrontal cortex. *Social Neuroscience*, *11*(4), 353 – 364.

Chib, V. S., Yun, K., Takahashi, H., & Shimojo, S. (2013). Noninvasive remote activation of the ventral midbrain by transcranial direct current stimulation of prefrontal cortex. *Translational Psychiatry*, *3*(6), e268.

Christoff, K., & Gabrieli, J. D. (2000). The frontopolar cortex and human cognition: Evidence for a rostrocaudal hierarchical organization within the human prefrontal cortex. *Psychobiology*, *28*(2), 168 – 186.

Chrysikou, E. G., Gorey, C., & Aupperle, R. L. (2017). Anodal transcranial direct current stimulation over right dorsolateral prefrontal cortex alters decision making during approach-avoidance conflict. *Social Cognitive and Affective Neuroscience*, *12*(3), 468 – 475.

Clarke, P. J. F., Browning, M., Hammond, G., Notebaert, L., & MacLeod, C. (2014). The causal role of the dorsolateral prefrontal cortex in the modification of attentional bias: Evidence from transcranial direct current stimulation. *Biological Psychiatry*, *76*(12), 946 – 952.

Cooper, J. A., Gorlick, M. A., Denny, T., Worthy, D. A., Beevers, C. G., & Todd Maddox, W. (2014). Training attention improves decision making in individuals with elevated self-reported depressive symptoms. *Cognitive, Affective and Behavioral Neuroscience, 14*(2), 729 – 741.

Creutzfeld, O. D., Fromm, G. H. & Kapp, H. (1962). Influence of transcortical dc-currents on cortical neuronal activity. *Experimental Neurology, 5*, 436 – 452.

Crowther, L. J., Marketos, P., Williams, P. I., Melikhov, Y., Jiles, D. C., & Starzewski, J. H. (2011). Transcranial magnetic stimulation: Improved coil design for deep brain investigation. *Journal of Applied Physics, 109*(7), 07B314.

Csifcsák, G., Boayue, N. M., Puonti, O., Thielscher, A., & Mittner, M. (2018). Effects of transcranial direct current stimulation for treating depression: A modeling study. *Journal of Affective Disorders, 234*, 164 – 173.

Dalgleish, T., & Watts, F. N. (1990). Biases of attention and memory in disorders of anxiety and depression. *Clinical Psychology Review, 10*(5), 589 – 604.

Datta, A., Bansal, V., Diaz, J., Patel, J., Reato, D., & Bikson, M. (2009). Gyri-precise head model of transcranial direct current stimulation: Improved spatial focality using a ring electrode versus conventional rectangular pad. *Brain Stimulation: Basic, Translational, and Clinical Research in Neuromodulation*, *2*(4), 201 – 207.

Davidson, R. J., Irwin, W., Anderle, M. J., & Kalin, N. H. (2003). The neural substrates of affective processing in depressed patients treated with venlafaxine. *American Journal of Psychiatry, 160*(1), 64 – 75.

Daw, N. D., Gershman, S. J., Seymour, B., Dayan, P., & Dolan, R. J. (2011). Model-based influences on humans' choices and striatal prediction errors. *Neuron, 69*(6), 1204 – 1215.

Dayan, P., & Niv, Y. (2008). Reinforcement learning: The Good, The Bad and The Ugly. *Current Opinion in Neurobiology, 18*(2), 185 – 196.

De Berker, A. O., Bikson, M., & Bestmann, S. (2013). Predicting the behavioral impact of transcranial direct current stimulation: Issues and limitations. *Frontiers in Human Neuroscience, 7*, 613.

Dedoncker, J., Brunoni, A. R., Baeken, C., & Vanderhasselt, M.-A. (2016). A systematic review and meta-analysis of the effects of transcranial direct current stimulation (tDCS) over the dorsolateral prefrontal cortex in healthy and neuropsychiatric samples: Influence of stimulation parameters. *Brain Stimulation*, *9*(4), 501 – 517.

Dell'Osso, B., Dobrea, C., Arici, C., Benatti, B., Ferrucci, R., Vergari, M., ... & Altamura, A. C. (2014). Augmentative transcranial direct current stimulation (tDCS) in poor responder depressed patients: A follow-up study. *CNS Spectrums*, *19*(4), 347 – 354.

Dell'Osso, B., Zanoni, S., Ferrucci, R., Vergari, M., Castellano, F., D'Urso, N., ... & Altamura, A. C. (2012). Transcranial direct current stimulation for the outpatient treatment of poor-responder depressed patients. *European Psychiatry*, *27*(7), 513 – 517.

DeRubeis, R. J., Gelfand, L. A., Tang, T. Z., & Simons, A. D. (1999). Medications versus cognitive behavior therapy for severely depressed outpatients: Mega-analysis of four randomized comparisons. *American Journal of Psychiatry*, *156*(7), 1007 – 1013.

De Voogd, E. L., Wiers, R. W., Prins, P. J. M., De Jong, P. J., Boendermaker, W. J., Zwitser, R. J., & Salemink, E. (2016). Online attentional bias modification training targeting anxiety and depression in unselected adolescents: Short-and long-term effects of a randomized controlled trial. *Behaviour Research and Therapy*, *87*, 11 – 22.

Dichter, G. S., Smoski, M. J., Kampov-Polevoy, A. B., Gallop, R., & Garbutt, J. C. (2010). Unipolar depression does not moderate responses to the Sweet Taste Test. *Depression and Anxiety*, *27*(9), 859 – 863.

Dickens, C., McGowan, L., & Dale, S. (2003). Impact of depression on experimental pain perception: A systematic review of the literature with meta-analysis. *Psychosomatic Medicine*, *65*(3), 369 – 375.

Disner, S. G., Beevers, C. G., Haigh, E. A. P., & Beck, A. T. (2011). Neural mechanisms of the cognitive model of depression. *Nature Reviews Neuroscience*, *12*(8), 467 – 477.

Dmochowski, J. P., Datta, A., Huang, Y., Richardson, J. D., Bikson, M., Fridriksson, J., & Parra, L. C. (2013). Targeted transcranial direct current stimulation for rehabilitation after stroke. *Neuroimage*, *75*, 12 – 19.

Doll, B. B., Simon, D. A., & Daw, N. D. (2012). The ubiquity of model-based reinforcement learning. *Current Opinion in Neurobiology*, *22*(6), 1075 – 1081.

Dombrovski, A. Y., Clark, L., Siegle, G. J., Butters, M. A, Ichikawa, N., Sahakian, B. J., & Szanto, K. (2010). Reward / punishment reversal learning in older suicide attempters. *American Journal of Psychiatry*, *167*, 699 – 707.

Dombrovski, A. Y., Szanto, K., Clark, L., Aizenstein, H. J., Chase, H. W., Reynolds 3rd, C. F., & Siegle, G. J. (2015). Corticostriatothalamic reward prediction error signals and executive control in late-life depression. *Psychological Medicine*, *45*(7), 1413 – 1424.

Domenech, P., & Koechlin, E. (2015). Executive control and decision-making in the prefrontal cortex. *Current Opinion in Behavioral Sciences*, *1*, 101 – 106.

Douglas, K. M., & Porter, R. J. (2010). Recognition of disgusted facial expressions in severe depression. *The British Journal of Psychiatry*, *197*(2), 156 – 157.

Doya, K. (2002). Metalearning and neuromodulation. *Neural Networks*, *15*(4–6), 495 – 506.

Doya, K. (2007). Reinforcement learning: Computational theory and biological mechanisms. *HFSP Journal*, *1*(1), 30 – 40.

Doya, K. (2008). Modulators of decision making. *Nature Neuroscience*, *11*(4), 410 – 416.

Dutta, A., McKie, S., & Deakin, J. W. (2014). Resting state networks in major depressive disorder. *Psychiatry Research: Neuroimaging*, *224*(3), 139 – 151.

Edwards, D., Cortes, M., Datta, A., Minhas, P., Wassermann, E. M., & Bikson, M. (2013). Physiological and modeling evidence for focal transcranial electrical brain stimulation in humans: a basis for high-definition tDCS. *NeuroImage*, *74*, 266 – 275.

Ekman, P. & Friesen, W. (1976). *Pictures of facial affect.* Palo Alto, CA: Consulting Psychologists Press.

Ernst, M., Nelson, E. E., McClure, E. B., Monk, C. S., Munson, S., Eshel, N., ... & Blair, J. (2004). Choice selection and reward anticipation: An fMRI study. *Neuropsychologia*, *42*(12), 1585 – 1597.

Eshel, N., & Roiser, J. P. (2010). Reward and punishment processing in depression. *Biological Psychiatry*, *68*(2), 118 – 124.

Faul, F., Erdtfelder, E., Lang, A.-G., & Buchner, A. (2007). G*Power 3: A flexible statistical power analysis programme for the social, behavioral, and biomedical sciences. *Behavior Research Methods*, *39*, 175 – 191.

Fecteau, S., Knoch, D., Fregni, F., Sultani, N., Boggio, P., & Pascual-Leone, A. (2007a). Diminishing risk-taking behavior by modulating activity in the prefrontal cortex: A direct current stimulation study. *Journal of Neuroscience*, *27*(46), 12500 – 12505.

Fecteau, S., Pascual-Leone, A., Zald, D. H., Liguori, P., Theoret, H., Boggio, P. S., & Fregni, F. (2007b). Activation of prefrontal cortex by transcranial direct current

stimulation reduces appetite for risk during ambiguous decision making. *Journal of Neuroscience*, *27*(23), 6212 – 6218.

Ferrucci, R., Bortolomasi, M., Vergari, M., Tadini, L., Salvoro, B., Giacopuzzi, M., ... & Priori, A. (2009). Transcranial direct current stimulation in severe, drug-resistant major depression. *Journal of Affective Disorders*, *118*(1), 215 – 219.

Fertonani, A., Ferrari, C., & Miniussi, C. (2015). What do you feel if I apply transcranial electric stimulation? Safety, sensations and secondary induced effects. *Clinical Neurophysiology*, *126*(11), 2181 – 2188.

Filmer, H. L., Dux, P. E., & Mattingley, J. B. (2014). Applications of transcranial direct current stimulation for understanding brain function. *Trends in Neurosciences*, *37*(12), 742 – 753.

First, M.B., Williams, J.B.W., Karg, R.S., & Spitzer, R.L. (2015). Structured Clinical Interview for DSM-5 Disorders, Clinician Version (SCID-5-CV). Arlington, VA: American Psychiatric Association.

Fitz, N. S., & Reiner, P. B. (2013). The challenge of crafting policy for do-it-yourself brain stimulation. *Journal of Medical Ethics*, *41*: 413 – 414.

Fitzgerald, P. B., Maller, J. J., Hoy, K. E., Thomson, R., & Daskalakis, Z. J. (2009). Exploring the optimal site for the localization of dorsolateral prefrontal cortex in brain stimulation experiments. *Brain Stimulation: Basic, Translational, and Clinical Research in Neuromodulation*, *2*(4), 234 – 237.

Fleiss, J.L. (1986). *The Design and Analysis of Clinical Experiments.* New York: John Wiley and Sons.

Fregni, F., Boggio, P. S., Nitsche, M. A., Marcolin, M. A., Rigonatti, S. P., & Pascual-Leone, A. (2006a). Treatment of major depression with transcranial direct current stimulation. *Bipolar Disorders*, *8*(2), 203 – 204.

Fregni, F., Boggio, P. S., Santos, M. C., Lima, M., Vieira, A. L., Rigonatti, S. P., ... & Pascual-Leone, A. (2006b). Noninvasive cortical stimulation with transcranial direct current stimulation in Parkinson's disease. *Movement Disorders*, *21*(10), 1693 – 1702.

Fritsch, B., Reis, J., Martinowich, K., Schambra, H. M., Ji, Y., Cohen, L. G., & Lu, B. (2010). Direct current stimulation promotes BDNF-dependent synaptic plasticity: Potential implications for motor learning. *Neuron*, *66*(2), 198 – 204.

Funahashi, S., & Andreau, J. M. (2013). Prefrontal cortex and neural mechanisms of executive function. *Journal of Physiology Paris*, *107*(6), 471 – 482.

Gandiga, P. C., Hummel, F. C., & Cohen, L. G. (2006). Transcranial DC stimulation (tDCS): A tool for double-blind sham-controlled clinical studies in brain stimulation. *Clinical Neurophysiology*, *117*(4), 845 – 850.

Garrett, N., Sharot, T., Faulkner, P., Korn, C. W., Roiser, J. P., & Dolan, R. J. (2014). Losing the rose tinted glasses: neural substrates of unbiased belief updating in depression. *Frontiers in Human Neuroscience*, *8*, 1 – 9.

Garrison, J., Erdeniz, B., & Done, J. (2013). Prediction error in reinforcement learning: A meta-analysis of neuroimaging studies. *Neuroscience and Biobehavioral Reviews*, *37*(7), 1297 – 1310.

Gaynes, B. N., Lloyd, S. W., Lux, L., Gartlehner, G., Hansen, R. A., Brode, S., ... & Lohr, K. N. (2014). Repetitive transcranial magnetic stimulation for treatment-resistant depression: A systematic review and meta-analysis. *The Journal of Clinical Psychiatry*, *75*(5), 477 – 489.

Geng, J. J., & Vossel, S. (2013). Re-evaluating the role of TPJ in attentional control: Contextual updating? *Neuroscience and Biobehavioral Reviews*, *37*(10), 2608 – 2620.

George, M. S., Lisanby, S. H., Avery, D., McDonald, W. M., Durkalski, V., Pavlicova, M., ... & Holtzheimer, P. E. (2010). Daily left prefrontal transcranial magnetic stimulation therapy for major depressive disorder: A sham-controlled randomized trial. *Archives of General Psychiatry, 67*(5), 507 – 516.

George, M. S., Taylor, J. J., & Short, E. B. (2013). The expanding evidence base for rTMS treatment of depression. *Current Opinion in Psychiatry, 26*(1), 13 – 18.

Geuzaine, C., & Remacle, J.-F (2009). Gmsh: A three-dimensional finite element mesh generator with built-in pre- and post-processing facilities. *International Journal for Numerical Methods in Engineering, 79*(11), 1309 – 1331.

Gläscher, J., Daw, N., Dayan, P., & O'Doherty, J. P. (2010). States versus rewards: dissociable neural prediction error signals underlying model-based and model-free reinforcement learning. *Neuron, 66*(4), 585 – 595.

Gläscher, J. P., & O'Doherty, J. P. (2010). Model-based approaches to neuroimaging: Combining reinforcement learning theory with fMRI data. *Wiley Interdisciplinary Reviews: Cognitive Science, 1*(4), 501 – 510.

Gollan, J. K., Pane, H. T., McCloskey, M. S., & Coccaro, E. F. (2008). Identifying differences in biased affective information processing in major depression. *Psychiatry Research, 159*(1), 18 – 24.

Gradin, V. B., Kumar, P., Waiter, G., Ahearn, T., Stickle, C., Milders, M., … Steele, J. D. (2011). Expected value and prediction error abnormalities in depression and schizophrenia. *Brain, 134*(6), 1751 – 1764.

Gray, J.A., & McNaughton, N. (2000). *The neuropsychology of anxiety: An enquiry into the functions of the septo-hippocampal system.* Oxford, England: Oxford University Press.

Greenberg, P. E., Fournier, A.A., Sisitsky, T., Pike, C.T., & Kessler, R.C. (2015). The economic burden of adults with major depressive disorder in the United States (2005 and 2010). *Journal of Clinical Psychiatry, 76*(2), 155 – 162.

Grimm, S., Beck, J., Schuepbach, D., Hell, D., Boesiger, P., Bermpohl, F., ... & Northoff, G. (2008). Imbalance between left and right dorsolateral prefrontal cortex in major depression is linked to negative emotional judgment: An fMRI study in severe major depressive disorder. *Biological Psychiatry*, *63*(4), 369 – 376.

Hakamata, Y., Lissek, S., Bar-Haim, Y., Britton, J. C., Fox, N. A., Leibenluft, E., ... & Pine, D. S. (2010). Attention bias modification treatment: A meta-analysis toward the establishment of novel treatment for anxiety. *Biological Psychiatry*, *68*(11), 982 – 990.

Hallett, M. (2000). Transcranial magnetic stimulation and the human brain. *Nature*, *406*(6792), 147 – 150.

Hallion, L. S., & Ruscio, A. M. (2011). A meta-analysis of the effect of cognitive bias modification on anxiety and depression. *Psychological Bulletin*, *137*(6), 940 – 958.

Hämmerer, D., Bonaiuto, J., Klein-Flügge, M., Bikson, M., & Bestmann, S. (2016). Selective alteration of human value decisions with medial frontal tDCS is predicted by changes in attractor dynamics. *Scientific Reports*, *6*, 1 – 13.

Harmer, C. J., Bhagwagar, Z., Perrett, D. I., Völlm, B. A., Cowen, P. J., & Goodwin, G. M. (2003). Acute SSRI administration affects the processing of social cues in healthy volunteers. *Neuropsychopharmacology*, *28*(1), 148 – 152.

Harmer, C. J., Thilo, K. V., Rothwell, J. C., & Goodwin, G. M. (2001). Transcranial magnetic stimulation of medial–frontal cortex impairs the processing of angry facial expressions. *Nature Neuroscience*, *4*(1), 17 – 18.

Harvill, L. (1991). Standard error of measurement. *Educational Measurement: Issues and Practice, 10(2),* 33 – 41.

He, Q., Chen, M., Chen, C., Xue, G., Feng, T., & Bechara, A. (2016). Anodal stimulation of the left DLPFC increases IGT scores and decreases delay discounting rate in healthy males. *Frontiers in Psychology, 7*, 1 – 9.

Hecht, D. (2010). Depression and the hyperactive right-hemisphere. *Neuroscience Research, 68*(2), 77 – 87.

Heeren, A., Baeken, C., Vanderhasselt, M. A., Philippot, P., & De Raedt, R. (2015). Impact of anodal and cathodal transcranial direct current stimulation over the left dorsolateral prefrontal cortex during attention bias modification: An eye-tracking study. *PLoS One, 10*(4), 1 – 16.

Heeren, A., Billieux, J., Philippot, P., De Raedt, R., Baeken, C., de Timary, P., … Vanderhasselt, M. A. (2017). Impact of transcranial direct current stimulation on attentional bias for threat: A proof-of-concept study among individuals with social anxiety disorder. *Social Cognitive and Affective Neuroscience, 12*(2), 251 – 260.

Hesse, S., Werner, C., Pohl, M., Rueckriem, S., Mehrholz, J., & Lingnau, M. L. (2005). Computerized arm training improves the motor control of the severely affected arm after stroke: A single-blinded randomized trial in two centers. *Stroke, 36*(9), 1960 – 1966.

Ho, K.-A., Bai, S., Martin, D., Alonzo, A., Dokos, S., Puras, P., & Loo, C. K. (2014). A pilot study of alternative transcranial direct current stimulation electrode montages for the treatment of major depression. *Journal of Affective Disorders, 167*, 251 – 258.

Holmes, A. J., & Pizzagalli, D. A. (2008). Response conflict and frontocingulate dysfunction in unmedicated participants with major depression. *Neuropsychologia*, *46*(12), 2904 – 2913.

Hone-Blanchet, A., Edden, R. A., & Fecteau, S. (2016). Online effects of transcranial direct current stimulation in real time on human prefrontal and striatal metabolites. *Biological Psychiatry*, *80*(6), 432 – 438.

Hsu, W. Y., Zanto, T. P., Anguera, J. A., Lin, Y. Y., & Gazzaley, A. (2015). Delayed enhancement of multitasking performance: Effects of anodal transcranial direct current stimulation on the prefrontal cortex. *Cortex*, *69*, 175 – 185.

Huys, Q. J. M., Daw, N. D., & Dayan, P. (2015). Depression: A decision-theoretic analysis. *Annual Review of Neuroscience*, *38*(1), 1 – 23.

Huys, Q. J. M., Maia, T. V., & Frank, M. J. (2016). Computational psychiatry as a bridge from neuroscience to clinical applications. *Nature Neuroscience*, *19*(3), 404 – 413.

Huys, Q. J., Pizzagalli, D. A., Bogdan, R., & Dayan, P. (2013). Mapping anhedonia onto reinforcement learning: A behavioural meta-analysis. *Biology of Mood & Anxiety Disorders*, *3*(1), 12.

Hyman, S. E., Malenka, R. C., & Nestler, E. J. (2006). Neural mechanisms of addiction: The role of reward-related learning and memory. *Annual Review of Neuroscience*, *29*, 565 – 598.

Ironside, M., O'Shea, J., Cowen, P. J., & Harmer, C. J. (2016). Frontal cortex stimulation reduces vigilance to threat: Implications for the treatment of depression and anxiety. *Biological Psychiatry*, *79*(10), 823 – 830.

IsHak, W. W., Mirocha, J., James, D., Tobia, G., Vilhauer, J., Fakhry, H., ... & Cohen, R. M. (2015). Quality of life in major depressive disorder before/after multiple steps

of treatment and one-year follow-up. *Acta Psychiatrica Scandinavica*, *131*(1), 51 – 60.

Ishii, S., Yoshida, W., & Yoshimoto, J. (2002). Control of exploitation-exploration meta-parameter in reinforcement learning. *Neural Networks*, *15*(4–6), 665 – 687.

Iuculano, T., & Kadosh, R. C. (2013). The mental cost of cognitive enhancement. *Journal of Neuroscience*, *33*(10), 4482 – 4486.

Jacobson, L., Koslowsky, M., & Lavidor, M. (2012). tDCS polarity effects in motor and cognitive domains: A meta-analytical review. *Experimental Brain Research*, *216*(1), 1 – 10.

Jaeggi, S. M., Buschkuehl, M., Shah, P., & Jonides, J. (2014). The role of individual differences in cognitive training and transfer. *Memory & Cognition*, *42*(3), 464 – 480.

Jasper, H.H. (1958). The ten-twenty electrode system of the International Federation. *Electroencephalography and Clinical Neurophysiology, 10*, 371 – 375.

Jessup, R. K., & O'Doherty, J. P. (2014). Distinguishing informational from value-related encoding of rewarding and punishing outcomes in the human brain. *European Journal of Neuroscience*, *39*(11), 2014 – 2026.

Jocham, G., Hunt, L. T., Near, J., & Behrens, T. E. J. (2014). A mechanism for value-guided choice based on the excitation-inhibition balance in prefrontal cortex. *Nature Neuroscience*, *15*(7), 960 – 961.

Jones, E. B., & Sharpe, L. (2017). Cognitive bias modification: A review of meta-analyses. *Journal of Affective Disorders, 223*, 175 – 183.

Kalu, U. G., Sexton, C. E., Loo, C. K., & Ebmeier, K. P. (2012). Transcranial direct current stimulation in the treatment of major depression: A meta-analysis. *Psychological Medicine*, *42*(9), 1791 – 1800.

Katz, B., Au, J., Buschkuehl, M., Abagis, T., Zabel, C., Jaeggi, S. M., & Jonides, J. (2017). Individual differences and long-term consequences of tDCS-augmented cognitive training. *Journal of Cognitive Neuroscience*, *29*(9), 1498 – 1508.

Keeser, D., Meindl, T., Bor, J., Palm, U., Pogarell, O., Mulert, C., ... & Padberg, F. (2011). Prefrontal transcranial direct current stimulation changes connectivity of resting-state networks during fMRI. *Journal of Neuroscience*, *31*(43), 15284 – 15293.

Kendler, K. S. (2008). Explanatory models for psychiatric illness. *American Journal of Psychiatry*, *165*(6), 695 – 702.

Kendler, K. S., Thornton, L. M., & Gardner, C. O. (2001). Genetic risk, number of previous depressive episodes, and stressful life events in predicting onset of major depression. *American Journal of Psychiatry*, *158*(4), 582 – 586.

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

Kim, D. Y., Lim, J. Y., Kang, E. K., You, D. S., Oh, M. K., Oh, B. M., & Paik, N. J. (2010). Effect of transcranial direct current stimulation on motor recovery in patients with subacute stroke. *American Journal of Physical Medicine & Rehabilitation*, *89*(11), 879 – 886.

Kim, S., Stephenson, M. C., Morris, P. G., & Jackson, S. R. (2014b). tDCS-induced alterations in GABA concentration within primary motor cortex predict motor learning and motor memory: A 7 T magnetic resonance spectroscopy study. *Neuroimage*, *99*, 237 – 243.

Kim, S. H., Yoon, H., Kim, H., & Hamann, S. (2015). Individual differences in sensitivity to reward and punishment and neural activity during reward and avoidance learning. *Social Cognitive and Affective Neuroscience*, *10*(9), 1219 – 1227.

Kincses, T. Z., Antal, A., Nitsche, M. A., Bártfai, O., & Paulus, W. (2003). Facilitation of probabilistic classification learning by transcranial direct current stimulation of the prefrontal cortex in the human. *Neuropsychologia*, *42*(1), 113 – 117.

Koenigs, M., & Grafman, J. (2009). The functional neuroanatomy of depression: Distinct roles for ventromedial and dorsolateral prefrontal cortex. *Behavioural Brain Research*, *201*(2), 239 – 243.

Koo, T. K., & Li, M. Y. (2016). A guideline of selecting and reporting intraclass correlation coefficients for reliability research. *Journal of Chiropractic Medicine*, *15*(2), 155 – 163.

Korn, C. W., Sharot, T., Walter, H., Heekeren, H. R., & Dolan, R. J. (2014). Depression is related to an absence of optimistically biased belief updating about future life events. *Psychological Medicine*, *44*(3), 579 – 592.

Kozel, F. A., & George, M. S. (2002). Meta-analysis of left prefrontal repetitive transcranial magnetic stimulation (rTMS) to treat depression. *Journal of Psychiatric Practice*, *8*(5), 270 – 275.

Krause, B., & Cohen Kadosh, R. (2014). Not all brains are created equal: the relevance of individual differences in responsiveness to transcranial electrical stimulation. *Frontiers in Systems Neuroscience*, *8*, 25.

Kronberg, G., Bridi, M., Abel, T., Bikson, M., & Parra, L. C. (2017). Direct current stimulation modulates LTP and LTD: activity dependence and dendritic effects. *Brain Stimulation: Basic, Translational, and Clinical Research in Neuromodulation*, *10*(1), 51 – 58.

Kruijt, A. W., Putman, P., & Van der Does, W. (2013). The effects of a visual search attentional bias modification paradigm on attentional bias in dysphoric

individuals. *Journal of Behavior Therapy and Experimental Psychiatry*, *44*(2), 248 – 254.

Kunisato, Y., Okamoto, Y., Ueda, K., Onoda, K., Okada, G., Yoshimura, S., … Yamawaki, S. (2012). Effects of depression on reward-based decision making and variability of action in probabilistic learning. *Journal of Behavior Therapy and Experimental Psychiatry*, *43*(4), 1088 – 1094.

Laskowski, C. S., Williams, R. J., Martens, K. M., Gruber, A. J., Fisher, K. G., & Euston, D. R. (2016). The role of the medial prefrontal cortex in updating reward value and avoiding perseveration. *Behavioural Brain Research*, *306*, 52 – 63.

Laureiro-Martínez, D., Canessa, N., Brusoni, S., Zollo, M., Hare, T., Alemanno, F., & Cappa, S. F. (2014). Frontopolar cortex and decision-making efficiency: Comparing brain activity of experts with different professional background during an exploration-exploitation task. *Frontiers in Human Neuroscience*, *7*, 1 – 10.

Lee, D., Rushworth, M. F. S., Walton, M. E., Watanabe, M., & Sakagami, M. (2007). Functional specialization of the primate frontal cortex during decision making. *Journal of Neuroscience*, *27*(31), 8170 – 8173.

Lejuez, C. W., Read, J. P., Kahler, C. W., Richards, J. B., Ramsey, S. E., Stuart, G. L., ... & Brown, R. A. (2002). Evaluation of a behavioral measure of risk taking: The Balloon Analogue Risk Task (BART). *Journal of Experimental Psychology: Applied*, *8*(2), 75 – 84.

Levasseur-Moreau, J., & Fecteau, S. (2012). Translational application of neuromodulation of decision-making. *Brain Stimulation*, *5*(2), 77 – 83.

Li, H., Wei, D., Browning, M., Du, X., Zhang, Q., & Qiu, J. (2016). Attentional bias modification (ABM) training induces spontaneous brain activity changes in young

women with subthreshold depression: A randomized controlled trial. *Psychological Medicine*, *46*(5), 909 – 920.

Li, L. M., Uehara, K., & Hanakawa, T. (2015). The contribution of interindividual factors to variability of response in transcranial direct current stimulation studies. *Frontiers in Cellular Neuroscience*, *9*, 181.

Liebetanz, D., Nitsche, M. A., Tergau, F., & Paulus, W. (2002). Pharmacological approach to the mechanisms of transcranial DC-stimulation-induced after-effects of human motor cortex excitability. *Brain*, *125*(10), 2238 – 2247.

Lo, A. H., Humphreys, M., Byrne, G. J., & Pachana, N. A. (2012). Test–retest reliability and practice effects of the Wechsler Memory Scale-III. *Journal of neuropsychology*, *6*(2), 212 – 231.

Loo, C. K., Alonzo, A., Martin, D., Mitchell, P. B., Galvez, V., & Sachdev, P. (2012). Transcranial direct current stimulation for depression: 3-Week, randomised, sham-controlled trial. *British Journal of Psychiatry*, *200*(1), 52 – 59.

Lüscher, C., & Malenka, R. C. (2012). NMDA receptor-dependent long-term potentiation and long-term depression (LTP/LTD). *Cold Spring Harbor Perspectives in Biology*, *4*(6), a005710.

MacKay, D. J. (2003). *Information theory, inference and learning algorithms*. Cambridge, UK: Cambridge University Press.

Maddox, W. T., Gorlick, M. A., Worthy, D. A., & Beevers, C. G. (2012). Depressive symptoms enhance loss-minimization, but attenuate gain-maximization in history-dependent decision-making. *Cognition*, *125*(1), 118 – 124.

Maeda, F., Paul, J., Maria, J., Topka, H., & Pascual-Leone, A. (2000). Modulation of corticospinal excitability by repetitive transcranial magnetic stimulation. *Clinical Neurophysiology*, *111*, 800 – 805.

Maia, T. V. (2009). Reinforcement learning, conditioning, and the brain: Successes and challenges. *Cognitive, Affective and Behavioral Neuroscience*, *9*(4), 343 – 364.

Maia, T. V., & Frank, M. J. (2011). From reinforcement learning models to psychiatric and neurological disorders. *Nature neuroscience*, *14*(2), 154 – 162.

Maroun, M. (2006). Stress reverses plasticity in the pathway projecting from the ventromedial prefrontal cortex to the basolateral amygdala. *European Journal of Neuroscience*, *24*(10), 2917 – 2922.

Marr, D. (1982). *Vision: A computational investigation into the human representation and processing of visual information.* New York: Freeman.

Martin, D. M., Liu, R., Alonzo, A., Green, M., Player, M. J., Sachdev, P., & Loo, C. K. (2013). Can transcranial direct current stimulation enhance outcomes from cognitive training? A randomized controlled trial in healthy participants. *International Journal of Neuropsychopharmacology*, *16*(9), 1927 – 1936.

Martin, D. M., Liu, R., Alonzo, A., Green, M., & Loo, C. K. (2014). Use of transcranial direct current stimulation (tDCS) to enhance cognitive training: Effect of timing of stimulation. *Experimental Brain Research, 232*(10), 3345 – 3351.

Mathews, A., & MacLeod, C. (2005). Cognitive vulnerability to emotional disorders. *Annual Review of Clinical Psychology*, *1*(1), 167 – 195.

Maxwell, S.C., Lau, M.Y., & Howard, G.S. (2015). Is psychology suffering from a replication crisis? What does "failure to replicate" really mean? *American Psychologist, 70*(6), 487 – 498.

Mayberg, H. S. (2003). Modulating dysfunctional limbic-cortical circuits in depression: Towards development of brain-based algorithms for diagnosis and optimised treatment. *British Medical Bulletin*, *65*(1), 193 – 207.

McCabe, C., Cowen, P. J., & Harmer, C. J. (2009). Neural representation of reward in recovered depressed patients. *Psychopharmacology*, *205*(4), 667 – 677.

McCabe, C., Woffindale, C., Harmer, C. J., & Cowen, P. J. (2012). Neural processing of reward and punishment in young people at increased familial risk of depression. *Biological Psychiatry*, *72*(7), 588 – 594.

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

McGuire, J. T., Nassar, M. R., Gold, J. I., & Kable, J. W. (2014). Functionally dissociable influences on learning rate in a dynamic environment. *Neuron*, *84*(4), 870 – 881.

Mengotti, P., Dombert, P. L., Fink, G. R., & Vossel, S. (2017). Disruption of the right temporoparietal junction impairs probabilistic belief updating. *The Journal of Neuroscience*, *37*(22), 5419 – 5428.

Meron, D., Hedger, N., Garner, M., & Baldwin, D. S. (2015). Transcranial direct current stimulation (tDCS) in the treatment of depression: Systematic review and meta-analysis of efficacy and tolerability. *Neuroscience & Biobehavioral Reviews*, *57*, 46 – 62.

Minarik, T., Berger, B., Althaus, L., Bader, V., Biebl, B., Brotzeller, F., ... & Leitner, M. (2016). The importance of sample size for reproducibility of tDCS effects. *Frontiers in Human Neuroscience*, *10*, 453.

Minhas, P., Bansal, V., Patel, J., Ho, J. S., Diaz, J., Datta, A., & Bikson, M. (2010). Electrodes for high-definition transcutaneous DC stimulation for applications in drug delivery and electrotherapy, including tDCS. *Journal of Neuroscience Methods*, *190*(2), 188 – 197.

Minichiello, L. (2009). TrkB signalling pathways in LTP and learning. *Nature Reviews Neuroscience*, *10*(12), 850 – 860.

Miranda, P. C., Lomarev, M., & Hallett, M. (2006). Modeling the current distribution during transcranial direct current stimulation. *Clinical Neurophysiology*, *117*(7), 1623 – 1629.

Mogoaşe, C., David, D., & Koster, E. H. (2014). Clinical efficacy of attentional bias modification procedures: An updated meta-analysis. *Journal of Clinical Psychology*, *70*(12), 1133 – 1157.

Moliadze, V., Antal, A., & Paulus, W. (2010). Electrode-distance dependent after-effects of transcranial direct and random noise stimulation with extracephalic reference electrodes. *Clinical Neurophysiology*, *121*(12), 2165 – 2171.

Monte-Silva, K., Kuo, M., Hessenthaler, S., Fresnoza, S., Liebetanz, D., Paulus, W., & Nitsche, M. A. (2013). Brain stimulation induction of late LTP-like plasticity in the human motor cortex by repeated non-invasive brain stimulation. *Brain Stimulation*, *6*(3), 424 – 432.

Morgan, H. M., Davis, N. J., & Bracewell, R. M. (2014). Does transcranial direct current stimulation to prefrontal cortex affect mood and emotional memory retrieval in healthy individuals? *PLoS One*, *9*(3), e92162.

Morris, R. G. M., Anderson, E., Lynch, G. A., & Baudry, M. (1986). Selective impairment of learning and blockade of long-term potentiation by an N-methyl-D-aspartate receptor antagonist, AP5. *Nature*, *319*(6056), 774 – 776.

Moser, E. I., Krobert, K. A., Moser, M. B., & Morris, R. G. (1998). Impaired spatial learning after saturation of long-term potentiation. *Science*, *281*(5385), 2038 – 2042.

Moutoussis, M., Shahar, N., Hauser, T. U., & Dolan, R. J. (2017). Computation in psychotherapy, or how computational psychiatry can aid learning-based psychological therapies. *Computational Psychiatry, 2*, 50 – 73.

Moutoussis, M., Story, G. W., & Dolan, R. J. (2015). The computational psychiatry of reward: Broken brains or misguided minds? *Frontiers in Psychology, 6*, 1 – 9.

Murphy, F., Michael, A., Robbins, T., & Sahakian, B. (2003). Neuropsychological impairment in patients with major depressive disorder: The effects of feedback on task performance. *Psychological Medicine, 33*, 455 – 467.

Must, A., Szabó, Z., Bódi, N., Szász, A., Janka, Z., & Kéri, S. (2006). Sensitivity to reward and punishment and the prefrontal cortex in major depression. *Journal of Affective Disorders*, *90*(2–3), 209 – 215.

Nejati, V., Salehinejad, M. A., & Nitsche, M. A. (2018). Interaction of the left dorsolateral prefrontal cortex (l-DLPFC) and right orbitofrontal cortex (OFC) in hot and cold executive functions: Evidence from transcranial direct current stimulation (tDCS). *Neuroscience*, *369*, 109 – 123.

NICE (2018, June 11). *Depression in adults.* Retrieved from https://www.nice.org.uk/guidance/qs8/chapter/introduction-and-overview.

Nieuwenhuis, S., Heslenfeld, D. J., von Geusau, N. J. A., Mars, R. B., Holroyd, C. B., & Yeung, N. (2005). Activity in human reward-sensitive brain areas is strongly context dependent. *NeuroImage*, *25*(4), 1302 – 1309.

Nihonsugi, T., Ihara, A., & Haruno, M. (2015). Selective increase of intention-based economic decisions by noninvasive brain stimulation to the dorsolateral prefrontal cortex. *Journal of Neuroscience*, *35*(8), 3412 – 3419.

Nitsche, M. A., Boggio, P. S., Fregni, F., & Pascual-Leone, A. (2009). Treatment of depression with transcranial direct current stimulation (tDCS): A review. *Experimental Neurology*, *219*(1), 14 – 19.

Nitsche, M. A., Fricke, K., Henschke, U., Schlitterlau, A., Liebetanz, D., Lang, N., ... & Paulus, W. (2003). Pharmacological modulation of cortical excitability shifts induced by transcranial direct current stimulation in humans. *The Journal of Physiology*, *553*(1), 293 – 301.

Nitsche, M. A., & Paulus, W. (2000). Excitability changes induced in the human motor cortex by weak transcranial direct current stimulation. *The Journal of Physiology*, *527*(3), 633 – 639.

Nitsche, M. A., Seeber, A., Frommann, K., Klein, C. C., Rochford, C., Nitsche, M. S., ... & Paulus, W. (2005). Modulating parameters of excitability during and after transcranial direct current stimulation of the human motor cortex. *The Journal of Physiology*, *568*(1), 291 – 303.

Norman, R. M., Davies, F., Malla, A. K., Cortese, L., & Nicholson, I. R. (1996). Relationship of obsessive—compulsive symptomatology to anxiety, depression and schizotypy in a clinical population. *British Journal of Clinical Psychology*, *35*(4), 553 – 566.

O'Connell, N. E., Cossar, J., Marston, L., Wand, B. M., Bunce, D., Moseley, G. L., & De Souza, L. H. (2012). Rethinking clinical trials of transcranial direct current stimulation: Participant and assessor blinding is inadequate at intensities of 2mA. *PloS One*, *7*(10), e47514.

O'Doherty, J. P. (2004). Reward representations and reward-related learning in the human brain: Insights from neuroimaging. *Current Opinion in Neurobiology*, *14*(6), 769 – 776.

O'Doherty, J. P., Hampton, A., & Kim, H. (2007). Model-based fMRI and its application to reward learning and decision making. *Annals of the New York Academy of Sciences*, *1104*, 35 – 53.

Oostenveld, R., & Praamstra, P. (2001). The five percent electrode system for high-resolution EEG and ERP measurements. *Clinical Neurophysiology*, *112*(4), 713 – 719.

Opitz, A., Paulus, W., Will, S., Antunes, A., & Thielscher, A. (2015). Determinants of the electric field during transcranial direct current stimulation. *NeuroImage*, *109*, 140 – 150.

O'Reardon, J. P., Solvason, H. B., Janicak, P. G., Sampson, S., Isenberg, K. E., Nahas, Z., ... & Demitrack, M. A. (2007). Efficacy and safety of transcranial magnetic stimulation in the acute treatment of major depression: A multisite randomized controlled trial. *Biological Psychiatry*, *62*(11), 1208 – 1216.

O'Reilly, J. X., Schuffelgen, U., Cuell, S. F., Behrens, T. E. J., Mars, R. B., & Rushworth, M. F. S. (2013). Dissociable effects of surprise and model update in parietal and anterior cingulate cortex. *Proceedings of the National Academy of Sciences*, *110*(38), E3660 – E3669.

O'Shea, J., Revol, P., Cousijn, H., Near, J., Petitet, P., Jacquin-Courtois, S., ... & Rossetti, Y. (2017). Induced sensorimotor cortex plasticity remediates chronic treatment-resistant visual neglect. *eLife*, *6*, e26602.

Ott, D. V. M., Ullsperger, M., Jocham, G., Neumann, J., & Klein, T. A. (2011). Continuous theta-burst stimulation (cTBS) over the lateral prefrontal cortex alters reinforcement learning bias. *NeuroImage*, *57*(2), 617 – 623.

Palm, U., Schiller, C., Fintescu, Z., Obermeier, M., Keeser, D., Reisinger, E., ... & Padberg, F. (2012). Transcranial direct current stimulation in treatment resistant

depression: A randomized double-blind, placebo-controlled study. *Brain Stimulation: Basic, Translational, and Clinical Research in Neuromodulation*, *5*(3), 242 – 251.

Palminteri, S., Clair, A. H., Mallet, L., & Pessiglione, M. (2012). Similar improvement of reward and punishment learning by serotonin reuptake inhibitors in obsessive-compulsive disorder. *Biological Psychiatry, 72*(3), 244 – 250.

Pandya, D. N., & Yeterian, E. H. (1991). Prefrontal cortex in relation to other cortical areas in rhesus monkey: architecture and connections. *Progress in Brain Research, 85*, 63 – 94.

Paus, T., Castro-Alamancos, M.A. & Petrides, M. (2001). Cortico-cortical connectivity of the human mid-dorsolateral frontal cortex and its modulation by repetitive transcranial magnetic stimulation. *European Journal of Neuroscience, 14*, 1405 – 1411.

Payzan-LeNestour, E., Dunne, S., Bossaerts, P., & O'Doherty, J. (2013). The neural representation of unexpected uncertainty during value-based decision making. *Neuron*, *79*(1), 191 – 201.

Peckham, A. D., McHugh, R. K., & Otto, M. W. (2010). A meta-analysis of the magnitude of biased attention in depression. *Depression and Anxiety*, *27*(12), 1135 – 1142.

Peña-Gómez, C., Sala-Lonch, R., Junqué, C., Clemente, I. C., Vidal, D., Bargalló, N., ... & Bartrés-Faz, D. (2012). Modulation of large-scale brain networks by transcranial direct current stimulation evidenced by resting-state functional MRI. *Brain Stimulation: Basic, Translational, and Clinical Research in Neuromodulation*, *5*(3), 252 – 263.

Perera, T., George, M. S., Grammer, G., Janicak, P. G., Pascual-Leone, A., & Wirecki, T. S. (2016). The clinical TMS society consensus review and treatment

recommendations for TMS therapy for major depressive disorder. *Brain Stimulation: Basic, Translational, and Clinical Research in Neuromodulation*, *9*(3), 336 – 346.

Petitet, P., O'Reilly, J. X., Goncalves, A. M., Salvan, P., Kitazawa, S., Johansen-Berg, H., & O'Shea, J. (2018). Causal explanation of individual differences in human sensorimotor memory formation. *bioRxiv*, 255091.

Petrides, M., & Pandya, D. N. (1999). Dorsolateral prefrontal cortex: comparative cytoarchitectonic analysis in the human and the macaque brain and corticocortical connection patterns. *European Journal of Neuroscience*, *11*(3), 1011 – 1036.

Petrides, M., & Pandya, D. N. (2002). Comparative cytoarchitectonic analysis of the human and the macaque ventrolateral prefrontal cortex and corticocortical connection patterns in the monkey. *European Journal of Neuroscience*, *16*(2), 291 – 310.

Pizzagalli, D. A., Jahn, A. L., & O'Shea, J. P. (2005). Toward an objective characterization of an anhedonic phenotype: A signal-detection approach. *Biological Psychiatry*, *57*(4), 319 – 327.

Pizzagalli, D. A., Iosifescu, D., Hallett, L. A., Ratner, K. G., & Fava, M. (2008). Reduced hedonic capacity in major depressive disorder: Evidence from a probabilistic reward task. *Journal of Psychiatric Research*, *43*(1), 76 – 87.

Plazier, M., Joos, K., Vanneste, S., Ost, J., & De Ridder, D. (2012). Bifrontal and bioccipital transcranial direct current stimulation (tDCS) does not induce mood changes in healthy volunteers: A placebo controlled study. *Brain Stimulation: Basic, Translational, and Clinical Research in Neuromodulation*, *5*(4), 454 – 461.

Poreisz, C., Boros, K., Antal, A., & Paulus, W. (2007). Safety aspects of transcranial direct current stimulation concerning healthy subjects and patients. *Brain Research Bulletin*, *72*(4–6), 208 – 214.

Pringle, A., Browning, M., Cowen, P. J., & Harmer, C. J. (2011). A cognitive neuropsychological model of antidepressant drug action. *Progress in Neuro-Psychopharmacology and Biological Psychiatry*, *35*(7), 1586 – 1592.

Priori, A., Berardelli, A., Rona, S., Accornero, N., & Manfredi, M. (1998). Polarization of the human motor cortex through the scalp. *Neuroreport*, *9*(10), 2257 – 2260.

Priori, A., Hallett, M., & Rothwell, J. C. (2009). Repetitive transcranial magnetic stimulation or transcranial direct current stimulation? *Brain Stimulation*, *2*(4), 241 – 245.

Pripfl, J., Neumann, R., Köhler, U., & Lamm, C. (2013). Effects of transcranial direct current stimulation on risky decision making are mediated by "hot" and "cold" decisions, personality, and hemisphere. *European Journal of Neuroscience*, *38*(12), 3778 – 3785.

Pulcu, E., & Browning, M. (2017). Affective bias as a rational response to the statistics of rewards and punishments. *eLife*, *6*, 1 – 15.

Purpura, D. P. & McMurtry, J. G. (1965). Intracellular activities and evoked potential changes during polarization of motor cortex. *Journal of Neurophysiology, 28*, 166 – 185.

Rahman, A., Reato, D., Arlotti, M., Gasca, F., Datta, A., Parra, L. C., & Bikson, M. (2013). Cellular effects of acute direct current stimulation: Somatic and synaptic terminal effects. *Journal of Physiology*, *591*(10), 2563 – 2578.

Raja Beharelle, A., Polania, R., Hare, T. A., & Ruff, C. C. (2015). Transcranial stimulation over frontopolar cortex elucidates the choice attributes and neural mechanisms

used to resolve exploration-exploitation trade-offs. *Journal of Neuroscience*, *35*(43), 14544 –14556.

Rampersad, S. M., Janssen, A. M., Lucka, F., Aydin, U., Lanfer, B., Lew, S., … Oostendorp, T. F. (2014). Simulating transcranial direct current stimulation with a detailed anisotropic human head model. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, *22*(3), 441 – 452.

Reinhart, R. M. G., Zhu, J., Park, S., & Woodman, G. F. (2015). Medial-frontal stimulation enhances learning in schizophrenia by restoring prediction error signaling. *Journal of Neuroscience*, *35*(35), 12232 – 12240.

Reis, J., Schambra, H. M., Cohen, L. G., Buch, E. R., Fritsch, B., Zarahn, E., ... & Krakauer, J. W. (2009). Noninvasive cortical stimulation enhances motor skill acquisition over multiple days through an effect on consolidation. *Proceedings of the National Academy of Sciences*, *106*(5), 1590 – 1595.

Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. *Classical Conditioning II: Current Research and Theory*, *21*(6), 64 – 99.

Ridderinkhof, K. R., Van Den Wildenberg, W. P. M., Segalowitz, S. J., & Carter, C. S. (2004). Neurocognitive mechanisms of cognitive control: The role of prefrontal cortex in action selection, response inhibition, performance monitoring, and reward-based learning. *Brain and Cognition*, *56*(2), 129 – 140.

Ridding, M. C., & Rothwell, J. C. (2007). Is there a future for therapeutic use of transcranial magnetic stimulation? *Nature Reviews Neuroscience*, *8*(7), 559 – 567.

Rioult-Pedotti, M. S., Friedman, D., & Donoghue, J. P. (2000). Learning-induced LTP in neocortex. *Science*, *290*(5491), 533 – 536.

Robinson, O. J., & Chase, H. W. (2017). Learning and choice in mood disorders: Searching for the computational parameters of anhedonia. *Computational Psychiatry, 1*, 208 – 233.

Roiser, J. P., Elliott, R., & Sahakian, B. J. (2012). Cognitive mechanisms of treatment in depression. *Neuropsychopharmacology*, *37*(1), 117 – 136.

Rothkirch, M., Tonn, J., Köhler, S., & Sterzer, P. (2017). Neural mechanisms of reinforcement learning in unmedicated patients with major depressive disorder. *Brain*, *140*(4), 1147 – 1157.

Roy, L. B., Sparing, R., Fink, G. R., & Hesse, M. D. (2015). Modulation of attention functions by anodal tDCS on right PPC. *Neuropsychologia, 74*, 96 – 107.

Ruffini, G., Wendling, F., Merlet, I., Molaee-Ardekani, B., Mekonnen, A., Salvador, R., ... & Miranda, P. C. (2013). Transcranial current brain stimulation (tCS): Models and technologies. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, *21*(3), 333 – 345.

Rush, A. J., Trivedi, M. H., Wisniewski, S. R., Nierenberg, A. A., Stewart, J. W., Warden, D., ... & McGrath, P. J. (2006). Acute and longer-term outcomes in depressed outpatients requiring one or several treatment steps: a STAR* D report. *American Journal of Psychiatry, 163*(11), 1905 – 1917.

Rushworth, M. F. S., Noonan, M. A. P., Boorman, E. D., Walton, M. E., & Behrens, T. E. (2011). Frontal cortex and reward-guided learning and decision-making. *Neuron*, *70*(6), 1054 – 1069.

Russell, M. J., Goodman, T., Pierson, R., Shepherd, S., Wang, Q., Groshong, B., & Wiley, D. F. (2013). Individual differences in transcranial electrical stimulation current density. *Journal of Biomedical Research*, *27*(6), 495 – 508.

Russo, R., Twyman, P., Cooper, N. R., Fitzgerald, P. B., & Wallace, D. (2017). When you can, scale up: Large-scale study shows no effect of tDCS in an ambiguous risk-taking task. *Neuropsychologia*, *104*, 133 – 143.

Rutledge, R. B., Moutoussis, M., Smittenaar, P., Zeidman, P., Taylor, T., Hrynkiewicz, L., … Dolan, R. J. (2017). Association of neural and emotional impacts of reward prediction errors with major depression. *JAMA Psychiatry*, *74*(8), 790 – 797.

Sadacca, B. F., Wikenheiser, A. M., & Schoenbaum, G. (2017). Toward a theoretical role for tonic norepinephrine in the orbitofrontal cortex in facilitating flexible learning. *Neuroscience*, *345*, 124 – 129.

Salvador, R., Mekonnen, A., Ruffini, G., & Miranda, P. C. (2010, August). Modeling the electric field induced in a high resolution realistic head model during transcranial current stimulation. In *Engineering in Medicine and Biology Society (EMBC), 2010 Annual International Conference of the IEEE* (pp. 2073-2076). IEEE.

Sandrini, M., Umiltà, C., & Rusconi, E. (2011). The use of transcranial magnetic stimulation in cognitive neuroscience: A new synthesis of methodological issues. *Neuroscience & Biobehavioral Reviews*, *35*(3), 516 – 536.

Santesso, D. L., Dillon, D. G., Birk, J. L., Holmes, A. J., Goetz, E., Bogdan, R., & Pizzagalli, D. A. (2008). Individual differences in reinforcement learning: Behavioral, electrophysiological, and neuroimaging correlates. *Neuroimage*, *42*(2), 807 – 816.

Scholl, J., & Klein-Flügge, M. (in press). Understanding psychiatric disease by capturing ecologically relevant features of learning and decision-making. *Behavioural Brain Research*, 1 – 20.

Scholl, J., Kolling, N., Nelissen, N., Stagg, C. J., Harmer, C. J., & Rushworth, M. F. S. (2017). Excitation and inhibition in anterior cingulate predict use of past experiences. *eLife*, *6*, e20365.

Segrave, R. A., Arnold, S., Hoy, K., & Fitzgerald, P. B. (2014). Concurrent cognitive control training augments the antidepressant efficacy of tDCS: A pilot study. *Brain Stimulation: Basic, Translational, and Clinical Research in Neuromodulation*, *7*(2), 325 – 331.

Seligman. (1972). Learned helplessness. *Annual Review of Medicine, 23*, 407 – 412.

Shen, B., Yin, Y., Wang, J., Zhou, X., McClure, S. M., & Li, J. (2016). High-definition tDCS alters impulsivity in a baseline-dependent manner. *Neuroimage, 143*, 343 – 352.

Shiozawa, P., Fregni, F., Benseñor, I. M., Lotufo, P. A., Berlim, M. T., Daskalakis, J. Z., ... & Brunoni, A. R. (2014). Transcranial direct current stimulation for major depression: An updated systematic review and meta-analysis. *International Journal of Neuropsychopharmacology*, *17*(9), 1443 – 1452.

Shrout, P. E. (1998). Measurement reliability and reliability in psychiatry. *Statistical Methods in Medical Research*, *7*(98), 301 – 317.

Siebner, H., & Rothwell, J. (2003). Transcranial magnetic stimulation: New insights into representational cortical plasticity. *Experimental Brain Research*, *148*(1), 1 – 16.

Siegle, G. J., Thompson, W., Carter, C. S., Steinhauer, S. R., & Thase, M. E. (2007). Increased amygdala and decreased dorsolateral prefrontal BOLD responses in unipolar depression: Related and independent features. *Biological Psychiatry*, *61*(2), 198 – 209.

Smittenaar, P., FitzGerald, T. H. B., Romei, V., Wright, N. D., & Dolan, R. J. (2013). Disruption of dorsolateral prefrontal cortex decreases model-based in favor of model- free control in humans. *Neuron*, *80*(4), 914 – 919.

Smittenaar, P., Prichard, G., FitzGerald, T. H. B., Diedrichsen, J., & Dolan, R. J. (2014). Transcranial direct current stimulation of right dorsolateral prefrontal cortex does not affect model-based or model-free reinforcement learning in humans. *PLoS One*, *9*(1).

Smoski, M. J., Lynch, T. R., Rosenthal, M. Z., Cheavens, J. S., Chapman, A. L., & Krishnan, R. R. (2008). Decision-making and risk aversion among depressive adults. *Journal of Behavior Therapy and Experimental Psychiatry*, *39*(4), 567 – 576.

Sparing, R., & Mottaghy, F. M. (2008). Noninvasive brain stimulation with transcranial magnetic or direct current stimulation (TMS/tDCS) - From insights into human memory to therapy of its dysfunction. *Methods*, *44*(4), 329 – 337.

Spielberger, C. D. (1983). *STAI State-trait Anxiety Inventory for Adults Form Y: Review Set; Manual, Test, Scoring Key*. Mind Garden.

Stagg, C. J., Best, J. G., Stephenson, M. C., O'Shea, J., Wylezinska, M., Kincses, Z. T., ... & Johansen-Berg, H. (2009). Polarity-sensitive modulation of cortical neurotransmitters by transcranial stimulation. *Journal of Neuroscience*, *29*(16), 5202 – 5206.

Stagg, C. J., & Nitsche, M. A. (2011). Physiological basis of transcranial direct current stimulation. *Neuroscientist*, *17*(1), 37 – 53.

Stagg, C. J., Lin, R. L., Mezue, M., Segerdahl, A., Kong, Y., Xie, J., & Tracey, I. (2013). Widespread modulation of cerebral perfusion induced during and after transcranial

direct current stimulation applied to the left dorsolateral prefrontal cortex. *Journal of Neuroscience, 33*(28), 11425 – 11431.

Strafella, A. P., Paus, T., Fraraccio, M., & Dagher, A. (2003). Striatal dopamine release induced by repetitive transcranial magnetic stimulation of the human motor cortex. *Brain*, *126*(12), 2609 – 2615.

Surguladze, S. A., Young, A. W., Senior, C., Brébion, G., Travis, M. J., & Phillips, M. L. (2004). Recognition accuracy and response bias to happy and sad facial expressions in patients with major depression. *Neuropsychology*, *18*(2), 212 – 218.

Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning: An introduction* (Vol. 1, No. 1). Cambridge: MIT press.

Taylor Tavares, J. V., Clark, L., Furey, M. L., Williams, G. B., Sahakian, B. J., & Drevets, W. C. (2008). Neural basis of abnormal response to negative feedback in unmedicated mood disorders. *NeuroImage, 42*(3), 1118 – 1126.

Tecchio, F., Zappasodi, F., Assenza, G., Tombini, M., Vollaro, S., Barbati, G., & Rossini, P. M. (2010). Anodal transcranial direct current stimulation enhances procedural consolidation. *Journal of Neurophysiology*, *104*(2), 1134 – 1140.

Tervo, D. G. R., Proskurin, M., Manakov, M., Kabra, M., Vollmer, A., Branson, K., & Karpova, A. Y. (2014). Behavioral variability through stochastic choice and its gating by anterior cingulate cortex. *Cell*, *159*(1), 21 – 32.

Thielscher, A., Antunes, A., & Saturnino, G. B. (2015, August). Field modeling for transcranial magnetic stimulation: a useful tool to understand the physiological effects of TMS? In *Engineering in Medicine and Biology Society (EMBC), 2015 37th Annual International Conference of the IEEE* (pp. 222-225). IEEE.

Thoern, H. A., Grueschow, M., Ehlert, U., Ruff, C. C., & Kleim, B. (2016). Attentional bias towards positive emotion predicts stress resilience. *PLoS One, 11*(3), 1 – 8.

Thomsen, K. R. (2015). Measuring anhedonia: Impaired ability to pursue, experience, and learn about reward. *Frontiers in Psychology*, *6*, 1 – 11.

Tokic, M. (2010, September). Adaptive ε-greedy exploration in reinforcement learning based on value differences. In *Annual Conference on Artificial Intelligence* (pp. 203 – 210). Berlin, Heidelberg: Springer.

Tranter, R., Bell, D., Gutting, P., Harmer, C., Healy, D., & Anderson, I. M. (2009). The effect of serotonergic and noradrenergic antidepressants on face emotion processing in depressed patients. *Journal of Affective Disorders*, *118*(1), 87 – 93.

Treadway, M. T., Bossaller, N., Shelton, R. C., & Zald, D. H. (2012). Effort-based decision-making in major depressive disorder: A translational model of motivational anhedonia. *Journal of Abnormal Psychology*, *121*(3), 553 – 558.

Trew, J. L. (2011). Exploring the roles of approach and avoidance in depression: An integrative model. *Clinical Psychology Review*, *31*(7), 1156 – 1168.

Tsuchida, A., Doll, B. B., & Fellows, L. K. (2010). Beyond reversal: A critical role for human orbitofrontal cortex in flexible learning from probabilistic feedback. *Journal of Neuroscience*, *30*(50), 16868 – 16875.

Turi, Z., Mittner, M., Opitz, A., Popkes, M., Paulus, W., & Antal, A. (2015). Transcranial direct current stimulation over the left prefrontal cortex increases randomness of choice in instrumental learning. *Cortex*, *63*, 145 – 154.

Ubl, B., Kuehner, C., Kirsch, P., Ruttorf, M., Diener, C., & Flor, H. (2015). Altered neural reward and loss processing and prediction error signalling in depression. *Social Cognitive and Affective Neuroscience*, *10*(8), 1102 – 1112.

Usher, M., Cohen, J. D., Servan-Schreiber, D., Rajkowski, J., & Aston-Jones, G. (1999). The role of locus coeruleus in the regulation of cognitive performance. *Science*, *283*, 549 – 554.

Valentin, V. V., Dickinson, A., & O'Doherty, J. P. (2007). Determining the neural substrates of goal-directed learning in the human brain. *Journal of Neuroscience*, *27*(15), 4019 – 4026.

Valiengo, L., Benseñor, I. M., Goulart, A. C., Oliveira, J. F., Zanao, T. A., Boggio, P. S., ... & Brunoni, A. R. (2013). The sertraline versus electrical current therapy for treating depression clinical study (select-TDCS): Results of the crossover and follow-up phases. *Depression and Anxiety*, *30*(7), 646 – 653.

Van Hees, M. L., Rotter, T., Ellermann, T., & Evers, S. M. (2013). The effectiveness of individual interpersonal psychotherapy as a treatment for major depressive disorder in adult outpatients: A systematic review. *BMC Psychiatry*, *13*(1), 22.

Vanderhasselt, M. A., De Raedt, R., Brunoni, A. R., Campanhã, C., Baeken, C., Remue, J., & Boggio, P. S. (2013). tDCS over the left prefrontal cortex enhances cognitive control for positive affective stimuli. *PLoS One*, *8*(5), 4 – 9.

Van Holstein, M., Froböse, M. I., O'Shea, J., Aarts, E., & Cools, R. (2018). Controlling striatal function via anterior frontal cortex stimulation. *Scientific Reports*, *8*(1), 3312.

Vickery, T. J., Chun, M. M., & Lee, D. (2011). Ubiquity and specificity of reinforcement signals throughout the human brain. *Neuron*, *72*(1), 166 – 177.

Walsh, J. J., & Han, M. H. (2014). The heterogeneity of ventral tegmental area neurons: Projection functions in a mood-related context. *Neuroscience*, *282*, 101 – 108.

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

Wan Lee, S., Shimojo, S., & O'Doherty, J. P. (2014). Neural computations underlying arbitration between model-based and model-free learning. *Neuron*, *81*(3), 687 – 699.

Warden, D., Rush, A. J., Trivedi, M. H., Fava, M., & Wisniewski, S. R. (2007). The STAR* D Project results: A comprehensive review of findings. *Current Psychiatry Reports*, *9*(6), 449 – 459.

Watson, D., Clark, L. A., & Tellegen, A. (1988). Development and validation of brief measures of positive and negative affect: The PANAS scales. *Journal of Personality and Social Psychology*, *54*(6), 1063 – 1070.

Weber, M. J., Messing, S. B., Rao, H., Detre, J. A., & Thompson-Schill, S. L. (2014). Prefrontal transcranial direct current stimulation alters activation and connectivity in cortical and subcortical reward systems: A tDCS-fMRI study. *Human Brain Mapping*, *35*(8), 3673 – 3686.

Weir, J. P. (2005). Quantifying test-retest reliability using the intraclass correlation coefficient and the SEM. *Journal of Strength and Conditioning Research*, *19*(1), 231 – 240.

Whiteson, S., Taylor, M. E., & Stone, P. (2007). Empirical studies in action selection with reinforcement learning. *Adaptive Behavior*, *15*(1), 33 – 50.

Whitton, A. E., Treadway, M. T., & Pizzagalli, D. A. (2015). Reward processing dysfunction in major depression bipolar disorder and schizophrenia. *Current Opinion in Psychiatry*, *28*(1), 7 – 12.

Wischnewski, M., Zerr, P., & Schutter, D. J. (2016). Effects of theta transcranial alternating current stimulation over the frontal cortex on reversal learning. *Brain Stimulation: Basic, Translational, and Clinical Research in Neuromodulation*, *9*(5), 705 – 711.

Wittkuhn, L., Eppinger, B., Bartsch, L. M., Thurm, F., Korb, F. M., & Li, S.-C. (2018). Repetitive transcranial magnetic stimulation over dorsolateral prefrontal cortex modulates value-based learning during sequential decision-making. *NeuroImage*, *167*, 384 – 395.

Woods, A. J., Antal, A., Bikson, M., Boggio, P. S., Brunoni, A. R., Celnik, P., … Nitsche, M. A. (2016). A technical guide to tDCS, and related non-invasive brain stimulation tools. *Clinical Neurophysiology*, *127*(2), 1031 – 1048.

World Health Organisation (2018, March 22). *Depression*. Retrieved from http://www.who.int/news-room/fact-sheets/detail/depression.

Wörsching, J., Padberg, F., Ertl-Wagner, B., Kumpf, U., Kirsch, B., & Keeser, D. (2016). Imaging transcranial direct current stimulation (tDCS) of the prefrontal cortex - correlation or causality in stimulation-mediated effects? *Neuroscience and Biobehavioral Reviews*, *69*, 333 – 356.

Wu, A. D., Fregni, F., Simon, D. K., Deblieck, C., & Pascual-Leone, A. (2008). Noninvasive brain stimulation for Parkinson's disease and dystonia. *Neurotherapeutics*, *5*(2), 345 – 361.

Wu, D., Qian, L., Zorowitz, R. D., Zhang, L., Qu, Y., & Yuan, Y. (2013). Effects on decreasing upper-limb poststroke muscle tone using transcranial direct current stimulation: A randomized sham-controlled study. *Archives of Physical Medicine and Rehabilitation*, *94*(1), 1 – 8.

Wu, X., Wang, T., Liu, C., Wu, T., Jiang, J., Zhou, D., & Zhou, J. (2017). Functions of learning rate in adaptive reward learning. *Frontiers in Human Neuroscience*, *11*, 1 – 15.

Xue, G., Xue, F., Droutman, V., Lu, Z.-L., Bechara, A., & Read, S. (2013). Common neural mechanisms underlying reversal learning by reward and punishment. *PLoS One*, *8*(12), e82169.

Yang, W., Ding, Z., Dai, T., Peng, F., & Zhang, J. X. (2015). Attention bias modification training in individuals with depressive symptoms: A randomized controlled trial. *Journal of Behavior Therapy and Experimental Psychiatry*, *49*, 101 – 111.

Yang, W., Zhang, J. X., Ding, Z., & Xiao, L. (2016). Attention bias modification treatment for adolescents with major depression: A randomized controlled trial. *Journal of the American Academy of Child & Adolescent Psychiatry*, *55*(3), 208 – 218.

Yarkoni, T., Poldrack, R. A., Nichols, T. E., Van Essen, D. C., & Wager, T. D. (2011). Large-scale automated synthesis of human functional neuroimaging data. *Nature Methods*, *8*, 665 – 670.

Ye, H., Chen, S., Huang, D., Wang, S., & Luo, J. (2015). Modulating activity in the prefrontal cortex changes decision-making for risky gains and losses: A transcranial direct current stimulation study. *Behavioural Brain Research*, *286*, 17 – 21.

Ye, H., Huang, D., Wang, S., Zheng, H., Luo, J., & Chen, S. (2016). Activation of the prefrontal cortex by unilateral transcranial direct current stimulation leads to an asymmetrical effect on risk preference in frames of gain and loss. *Brain Research*, *1648*, 325 – 332.

Yu, A. J., & Dayan, P. (2005). Uncertainty, neuromodulation, and attention. *Neuron*, *46*(4), 681 – 692.

Zajkowski, W. K., Kossut, M., & Wilson, R. C. (2017). A causal role for right frontopolar cortex in directed, but not random, exploration. *eLife*, *6*, 1 – 18.

Zangen, A., Roth, Y., Voller, B., & Hallett, M. (2005). Transcranial magnetic stimulation of deep brain regions: Evidence for efficacy of the H-coil. *Clinical Neurophysiology*, *116*(4), 775 – 779.

Zheng, X., Alsop, D. C., & Schlaug, G. (2011). Effects of transcranial direct current stimulation (tDCS) on human regional cerebral blood flow. *Neuroimage*, *58*(1), 26 – 33.

Ziemann, U., Iliać, T. V., Pauli, C., Meintzschel, F., & Ruge, D. (2004). Learning modifies subsequent induction of long-term potentiation-like and long-term depression-like plasticity in human motor cortex. *Journal of Neuroscience*, *24*(7), 1666 – 1672.

Zuroff, D. C., Koestner, R., Moskowitz, D. S., McBride, C., Marshall, M., & Bagby, M. R. (2007). Autonomous motivation for therapy: A new common factor in brief treatments for depression. *Psychotherapy Research*, *17*(2), 137 – 147.

Appendix A. Comparator models

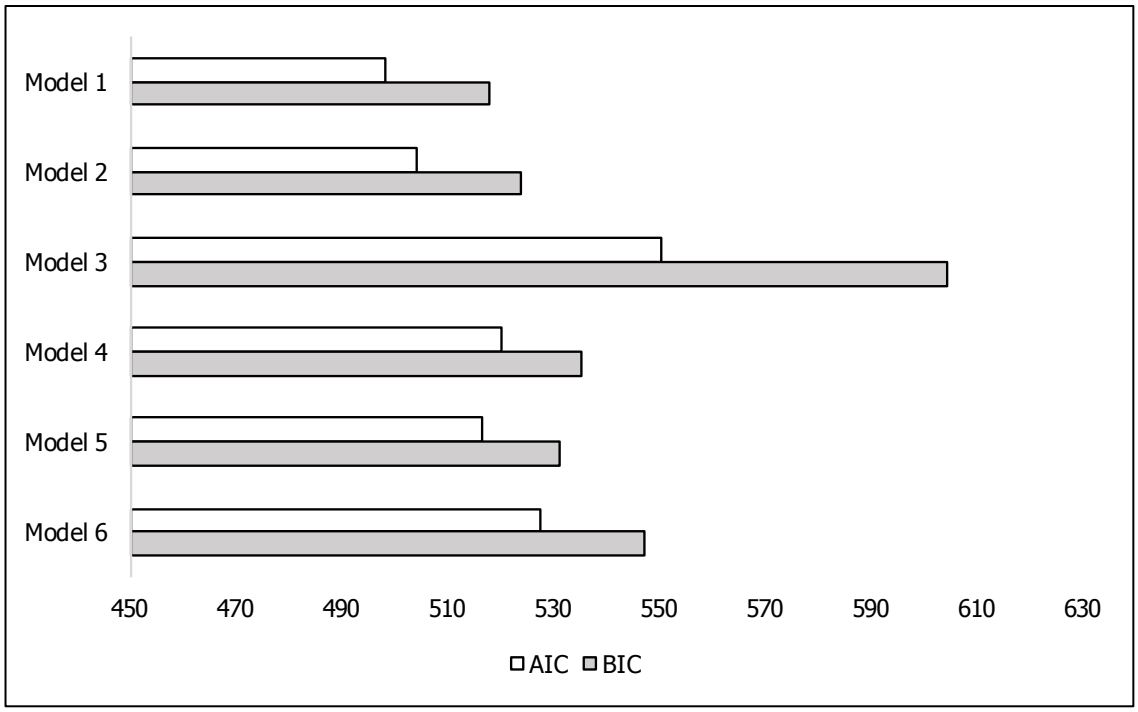

Figure A.1. AIC and BIC model fits for comparator models of the test-retest study (Chapter 3).

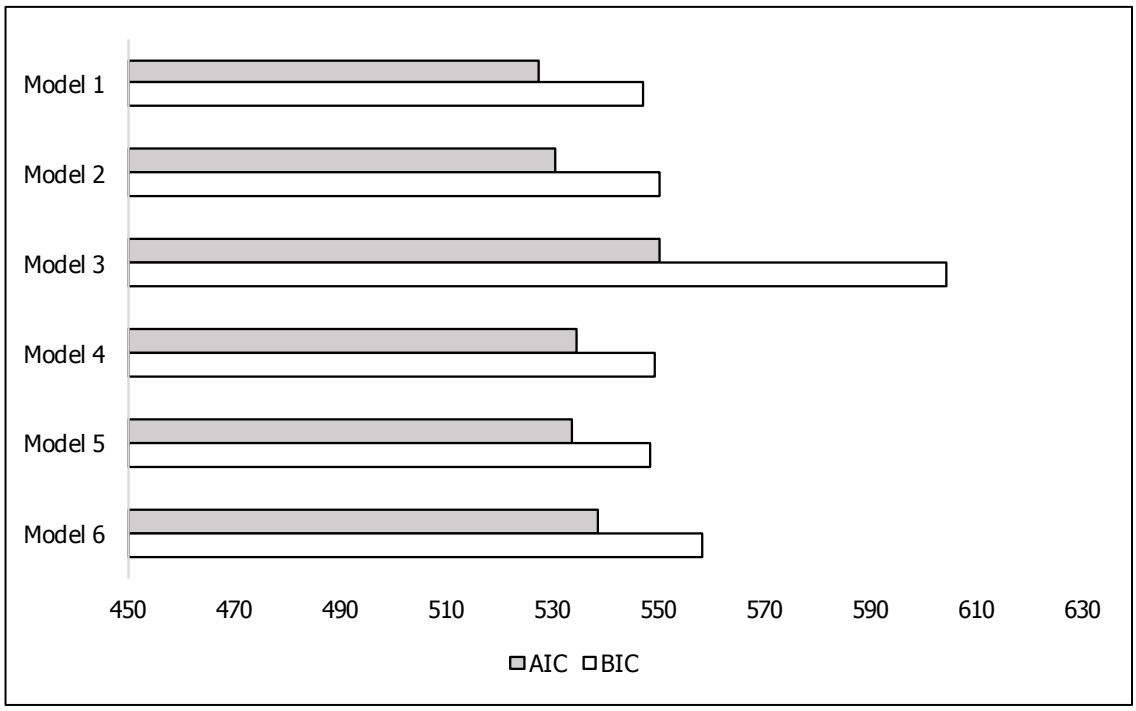

Figure A.2. AIC and BIC model fits for comparator models of the proof-of-concept tDCS study (Chapter 4).

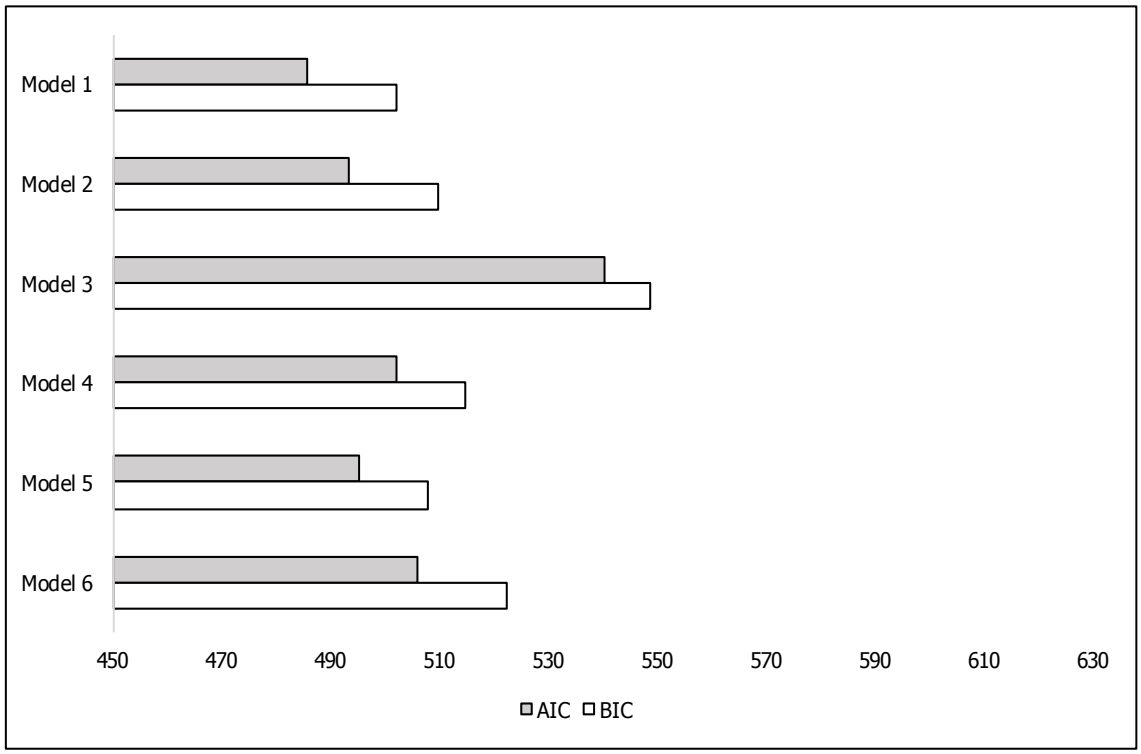

Figure A.3. AIC and BIC model fits for comparator models of the training tDCS study (Chapter 5).

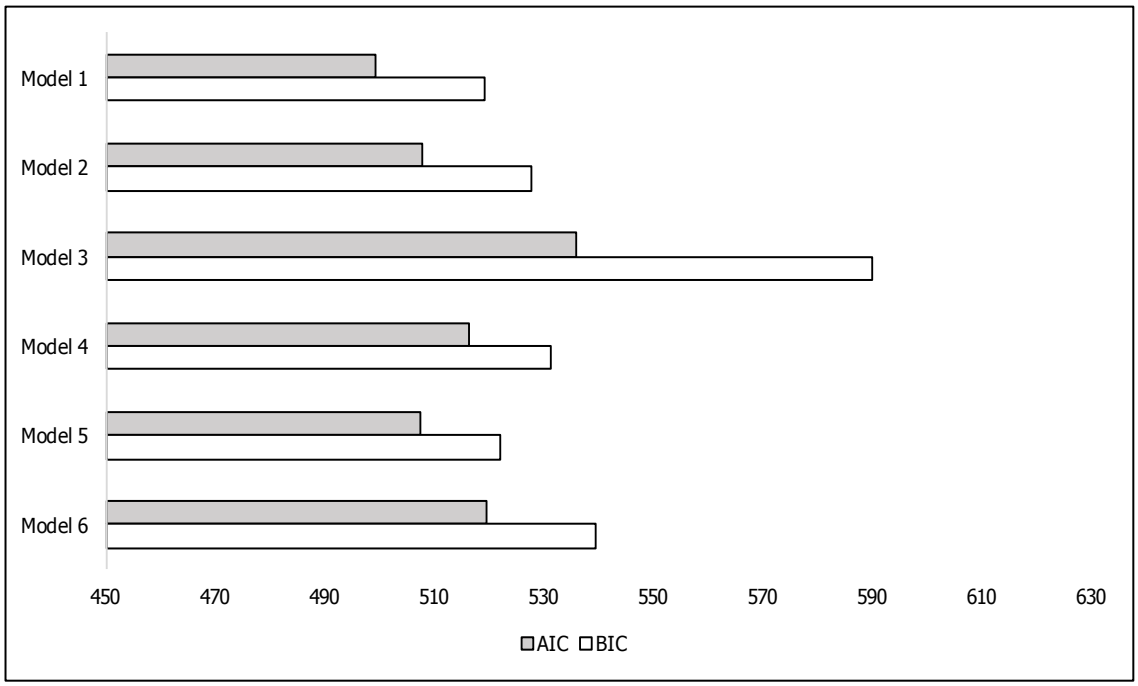

Figure A.4. AIC and BIC model fits for comparator models of the tDCS study on subclinical depression (Chapter 6).

Appendix B. Mood and anxiety correlations

Introduction

An exploratory aim of this research was to assess whether IBLT performance was correlated with trait measures of anxiety and depression. Previously, it has been reported that individuals with high anxiety tend to adapt their learning rates to a smaller extent in response to volatility changes compared with individuals with low anxiety scores on a punishment-based version of the Information Bias Learning Task (Browning et al., 2015). As suggested in Chapter 1, it can be hypothesised that depression may be associated with similar adaptation impairments. I therefore examined whether relatively high anxiety or low mood measures correlated with a decrease in learning rate differences between volatile and stable negative outcomes in Chapters 3, 4, and 6 (note that Chapter 5 was not included as participants completed either Win-volatile or Loss-volatile blocks, which did not allow for calculating volatile-stable contrasts). Win learning rates in the Loss-volatile blocks were subtracted from the Win-volatile blocks, and loss learning rates in the Win-volatile blocks were subtracted from the Loss-volatile blocks. For the tDCS studies, only data from the sham stimulation condition were investigated. Participants' average learning rates for the volatile-stable contrast were then correlated with BDI and STAI-Trait scores.

Correlations Chapter 3

Questionnaire data for mood and anxiety measures were missing for three participants, who were therefore excluded from the analyses. Correlation analyses were carried out across the blocks for both the BDI scores (M = 7.19, SE = 1.53) and the STAI-Trait scores (M = 39.48, SE = 1.94). The difference between volatile and stable learning rates for either win ($r(25) = 0.041$, $p = 0.839$) or loss outcomes ($r(25) = -0.013$, $p = 0.951$) was not

correlated with BDI scores. Similarly, STAI-Trait scores did not correlate with volatile-stable learning rates for either win ($r(25) = -0.531$, $p = 0.600$) or loss outcomes ($r(25) = -0.86$, $p = 0.398$).

Correlations Chapter 4

Online tDCS

For the online tDCS group, BDI and STAI-Trait data were unavailable for one participant, who was therefore excluded from the analyses. The volatility effect for losses (Loss-volatile – Win-volatile learning rates for loss outcomes) was not correlated with either BDI scores ($r(18) = -0.24$, $p = 0.148$) or STAI-Trait scores ($r(18) = 0.04$, $p = 0.82$). Similarly, there were no correlations between the volatility effect for wins (Win-volatile – Loss-volatile learning rates for win outcomes) and BDI ($r(18) = 0.29$, $p = 0.075$) or STAI-Trait scores ($r(18) = 0.23$, $p = 0.161$).

Offline tDCS

No significant associations were found between either the volatility effect for losses ($r(19) = 0.07$, $p = 0.651$) or wins ($r(19) = 0.21$, $p = 0.194$) on the BDI. In addition, there were no significant correlations between STAI-Trait scores and the volatility effect for either win outcomes ($r(19) = 0.19$, $p = 0.323$) or loss outcomes ($r(19) = 0.25$, $p = 0.125$).

Replication study

There was no significant correlation between loss volatility effects and BDI scores ($r(19) = 0.05$, $p = 0.779$) or STAI-Trait scores ($r(19) = -0.11$, $p = 0.497$). In addition, there were no correlations between either win volatility effects and the BDI ($r(19) = -0.30$, $p = 0.059$) or STAI-Trait values ($r(19) = -0.10$, $p = 0.532$).

Correlations Chapter 6

No significant correlations were observed for the loss volatility effect and either BDI scores ($r(8) = -0.26$, $p = 0.465$) or STAI-Trait scores ($r(8) = -0.17$, $p = 0.640$). Similarly, there were no correlations between the wins volatility effect and BDI scores ($r(8) = -0.16$, $p = 0.667$) or STAI-Trait scores ($r(8) = 0.27$, $p = 0.451$).

Comments

In contrast with previous findings (Browning et al., 2015), higher trait anxiety was not correlated with a smaller difference between learning rates for stable and changeable loss outcomes. The absence of a correlation may have been caused by the small range of anxiety scores obtained within these study samples. Whereas Browning and colleagues (2015) set out to recruit even numbers of participants with scores between the ranges of 20-30, 30-40, 40-50, and >50 on the STAI-Trait subscale, the majority of the participants in the present sample had a score between 30-50, with few participants demonstrating scores in the lowest or highest ranges. In addition, the previous study applied electric shocks rather than using monetary rewards and punishments, which may have resulted in different learning responses. I also did not find any correlations between indications of depressive symptomatology as indicated by BDI scores and adaptation to volatility for either wins or losses. Again, this lack of a correlation may have been due to a small score range, with 83% of the participants in Chapters 3 and 4 having a score of 10 or lower.

Appendix C. Model-based and model-free analyses

Introduction

There is preliminary evidence suggesting that stimulation of the DLPFC might change the ability to engage in model-based learning (Smittenaar et al., 2013; Wittkuhn et al., 2018). In order to assess potential changes in learning strategies due to stimulation, both model-based (Model 1) and model-free learning (Model 3) models as described in Chapter 3 were fitted to the participant data from Chapters 4-6. Akaike Information Criterion (AIC) and Bayesian Information Criterion (BIC) fits of these models were compared for sham and bifrontal tDCS conditions with repeated-measures ANOVAs using tDCS condition, Model type, Session, Run/Time, Outcome valence, and tDCS order as predictor variables. Lower AIC and BIC scores indicate a better fit of the model to the data. No significant alterations in either model-based or model-free learning strategies as a result of stimulation were observed across the tDCS studies reported here.

Chapter 4

Online tDCS

A repeated-measures ANOVA of the model fits showed that there was no interaction of Model type and tDCS condition for AIC values ($F(1,19) = 0.00$, $p = 0.955$) or BIC values ($F(1,19) = 0.00$, $p = 0.955$) in the Win- and Loss-volatile blocks. This finding indicates that participants' strategies when performing the IBLT were not altered by the application of bifrontal tDCS. Similarly, no effects of tDCS on learning strategy were found in the Both-volatile blocks (all $p > 0.05$).

Offline tDCS

There was a main effect of Model for BIC values ($F(1,19) = 46.95$, $p < 0.001$, $\eta^2 = 0.006$), with the model-based computations providing a better fit (M = 46.95, SE = 2.24) than the model-free approach (M = 50.30, SE = 2.19) in the Win- and Loss-volatile blocks. However, no such main effect was found for these models using AIC values ($F(1,19) = 1.40$, $p = 0.252$). In addition, similar to the online tDCS study, there was no tDCS Condition x Model type interaction for AIC or BIC values ($F(1,19) = 0.06$, $p = 0.807$), indicating that stimulation did not affect learning strategies in the offline tDCS study sample. There were no effects of either Model type or tDCS condition on AIC and BIC values in the Both-volatile blocks (all $p > 0.05$).

Replication study

For the BIC values of model-based versus model-free learning, there was again a main effect of Model type ($F(1,18) = 28.69$, $p < 0.001$, $\eta^2 = 0.004$), with the model-based computations providing a better fit (M = 52.80, SE = 2.21) than the model-free approach (M = 56.24, SE = 2.08). No such main effect was found for these models using AIC values ($F(1,18) = 0.84$, $p = 0.371$). There was no tDCS condition x Model type interaction for either AIC values ($F(1,18) = 0.00$, $p = 0.954$) or BIC values ($F(1,18) = 0.00$, $p = 0.955$), indicating that stimulation did not affect learning strategies. Furthermore, no effects of tDCS condition or Model type were found for the model fit in the Both-volatile blocks (all $p > 0.05$).

Chapter 5

Win-volatile training

A main effect was found for Model type in the Training blocks ($F(1,18) = 19.16$, $p < 0.001$, $\eta^2 = 0.010$), with Model-based learning providing a better fit (M= 49.19, SE = 3.01) than Model-free learning strategies (M = 50.44, SE = 2.86). A main effect of Model type was also observed in the Both-volatile blocks ($F(1,18) = 19.75$, $p < 0.001$, $\eta^2 = 0.064$). In addition, there was a Time x Model type interaction in the Both-volatile blocks ($F(1,18) = 4.48$, $p < 0.05$, $\eta^2 = 0.004$), which suggested that model fit for Model-based strategies decreased across the two Both-volatile blocks, whereas there was no change in Model-free learning fit (see Figure C.1). However, the interaction did not survive further exploration with paired *t*-tests (all $p > 0.05$). Resulted were unchanged if BIC values were investigated in the place of AIC fits.

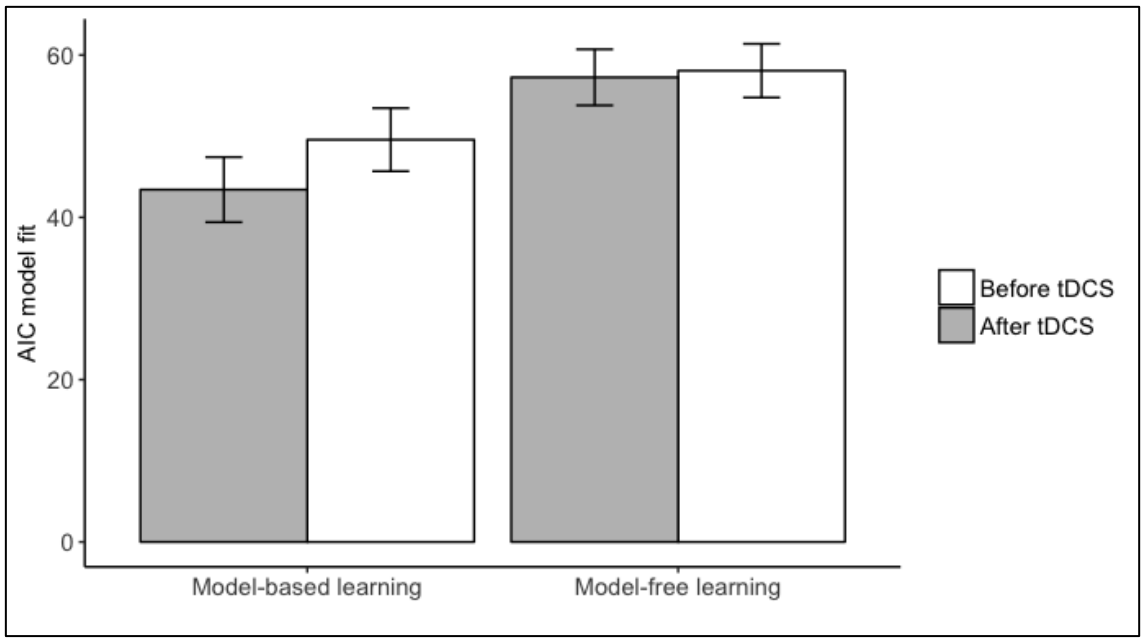

Figure C.1. Akaike Information Criterion (AIC) model fits for model-based and model-free learning strategies before and after tDCS and IBLT Win-volatile training.

Loss-volatile training

There were no effects of stimulation on the application of either Model-based or Model-free learning strategies in the AIC values for Training blocks (all p >0.05). In the Both-volatile blocks, a significant main effect of Model type was observed (F(1,18) = 9.47, p <0.01, η^2 = 0.019), with Model-based learning providing a closer fit to the data (M = 49.05, SE = 3.44) than Model-free learning (M = 56.75, SE = 2.88). Results were unchanged if BIC values were used instead of AIC values.

Chapter 6

In the Win- and Loss-volatile blocks, there was a main effect of Model type on AIC values (F(1,8) = 12.85, p <0.01, η^2 = 0.002), with Model-based learning (Model 1) providing a better fit to the data (M = 43.71, SE = 2.58) compared to Model-free (Model 3) learning (M = 46.13, SE = 2.40). In the Both-volatile blocks, a similar main effect of Model type was observed (F(1,8) = 11.09, p <0.05, η^2 = 0.024), with AIC values being smaller for Model 1 (M = 37.38, SE = 3.76) than for Model 3 (M = 41.66, SE = 3.78). In addition, a main effect of tDCS condition was found (F(1,8) = 11.09, p <0.05, η^2 = 0.024), with the AIC values for Model 1 (M = 43.10, SE = 3.37) and Model 3 (M = 45.57, SE = 3.35) being lower in the sham tDCS session than the AIC values for Model 1 (M = 44.32, SE = 3.94) and Model 3 (M = 46.79, SE = 3.49) in the active tDCS session. This indicates that individuals' behaviour was less well described by both models when receiving active tDCS compared to sham tDCS. However, as there were no interactions of tDCS condition with either Model type and/or Time (all p >0.05), there were no indications of a shift in the use of model-based versus model-free learning as a result of stimulation. Results were unchanged if BIC values rather than AIC values were used in the analyses.

Appendix D. tDCS safety screening form

Volunteer name: _____ Date of birth: _____

Have you ever suffered from any neurological or psychiatric conditions? If YES please give details (nature of condition, duration, current medication, etc).	***YES***	***NO***
Have you ever suffered from epilepsy, febrile convulsions in infancy or had recurrent fainting spells?	***YES***	***NO***
Does anyone in your immediate or distant family suffer from epilepsy? If YES please state your relationship to the affected family member.	***YES***	***NO***
Have you ever undergone a neurosurgical procedure (including eye surgery)? If YES please give details	***YES***	***NO***
Do you currently have any of the following fitted to your body? (please circle) ***Heart pacemaker Cochlear implant Medication pump Surgical clips***	***YES***	***NO***
Are you currently taking any unprescribed or prescribed medication? If YES please give details.	***YES***	***NO***
Are you currently undergoing anti - malarial treatment, or have been in the last 3 days?	***YES***	***NO***
Or drunk more than 3 units of alcohol in the last 24 hours?	***YES***	***NO***
Have you drunk alcohol already today?	***YES***	***NO***
Have you had more than one cup of coffee, or other sources of caffeine in the last hour?	***YES***	***NO***
Have you used recreational drugs in the last 24 hours?	***YES***	***NO***
Did you have very little sleep last night?	***YES***	***NO***
Is there any chance you might be pregnant?	***YES***	***NO***
Have you already participated in a tDCS experiment in the last week?	***YES***	***NO***

Volunteer signature: _____ Date of study: _____

Staff signature: _____ Scanner: _____

Appendix E. Adverse effects questionnaire

Participant number: _____

Date: _____

Session: _____

Code: _____

Did you experience any of the following during the session? (Please tick).					If **yes,** how strongly do you think it was related to tDCS? (Please tick).				
	Absent	Mild	Moderate	Severe	No	Remotely	Possibly	Probably	Definitely
Headache									
Neck pain									
Scalp pain									
Tingling									
Itching									
Burning									
Skin redness									
Sleepiness									
Fatigue									
Mood change									
Trouble concentrating									
Other									

Please indicate the level of **DISTRACTION** caused by the stimulation.

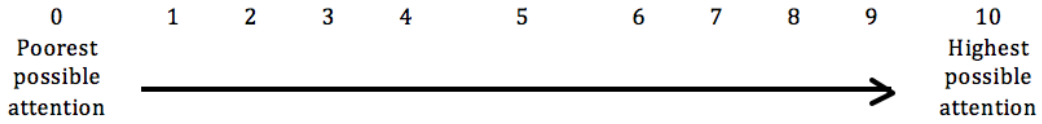

Do you think that you experienced real or sham tDCS during this session?

Real Sham

Based on adverse effect questionnaire proposed by Brunoni, Amadera, Berbel, Volz, Gomes Rizzerio and Felipe Fregni 2011. Modified by Stephanie Larcombe 22/09/14, Pierre Petitet 26/09/2014, and Margot Overman 13/01/2016

Appendix F. IBLT questionnaire

1. Did you use a particular strategy during the task?

2. Did you think that one of the outcomes (either wins or losses) was more useful than the other when doing the task?

AFTER SECOND SESSION ONLY:

3. Did you notice that one outcome was more changeable than the other?