UNIVERSITY OF
OXFORD

Cognitive approaches to non-invasive brain stimulation in the treatment of depression

Verena Sarrazin

Balliol College

Department of Psychiatry

University of Oxford

A thesis submitted for the degree of Doctor of Philosophy
under the supervision of
Prof Jacinta O'Shea and Prof Michael Browning.

Michaelmas Term 2022

Abstract

This thesis investigated two cognitive approaches for improving the application of non-invasive brain stimulation in the treatment of depression. The goal of the first project was to test whether combining transcranial direct current stimulation (tDCS) with a reinforcement learning task can reduce negative biases in depression. Recent research suggests that negative biases might arise from overestimating the informativeness of negative vs. positive information which might lead to increased learning from negative information. Two samples were recruited to test whether bifrontal tDCS applied during reinforcement learning increases reward vs. punishment learning in healthy participants and individuals with low mood. For the experimental paradigm, we used the Information Bias Learning Task which manipulates the relative informativeness of reward and punishments. As hypothesised, real compared to sham tDCS selectively increased reward learning in healthy individuals. To investigate whether this effect might be helpful to normalise reinforcement learning in depression, we compared performance between individuals with low mood and healthy participants. We found that individuals with low mood adjusted their punishment learning rate less, and their reward learning rate more to changes in informativeness than healthy individuals. Bifrontal tDCS normalised learning rate adjustment in individuals with low mood, which might be beneficial in depression treatment. The reported findings were only observed in response to bifrontal tDCS applied *during* (as opposed to *before*) task performance. This suggests the hypothesis that tDCS combined with reinforcement learning might have larger antidepressant effects than tDCS alone, which can be tested in future clinical trials. The second project was motivated by the idea of investigating different cognitive phenotypes of depression to guide treatment selection for individual patients. Previous research indicates that the relative severity of anxiety vs. anhedonia might differentiate between phenotypes that differ in the TMS target region they best respond to. We therefore recruited participants with varying symptom levels on an online platform to test which parameters extracted from five cognitive tasks differentially relate to symptoms of anxiety and anhedonia. The cognitive process which seemed to best differentiate between anxiety and anhedonia was punishment learning with volatile and stable

associations. Future clinical trials are needed to test whether these cognitive task parameters have any clinical value, for example for selecting the optimal TMS target region.

This thesis contains approximately 48,900 words.

Acknowledgements

When someone asks me how my DPhil went, I usually say, different from what I expected. Never would I have thought that my DPhil might be interrupted by a pandemic, causing a delay of 10 months. Two years ago, when I was writing yet another ethics amendment to get approval for restarting in-person testing, it seemed impossible that I would submit a thesis within the given timeframe. Despite these rather difficult conditions, I have learned a lot during my DPhil and enjoyed every stage of it. However, my DPhil would not have turned into such a valuable and enjoyable experience without the support and contributions from many wonderful people.

First of all, this work would not have been possible without my supervisors Jacinta and Mike who gave me helpful advice and support throughout all stages of my DPhil. I would like to thank you for encouraging me to explore my dataset in every single detail and try out different analysis approaches. Without your patience and our very detailed discussions, I would not have obtained the interesting findings reported in this thesis.

I would like to thank my fellow DPhil students and previous researchers in the TiNSL lab, Karen, Bronwyn, Emile and Margot, for all their scientific and social support. Your ideas and feedback have been incredibly helpful, and I cannot thank you enough for helping out in so many of my testing sessions. I am looking forward to having my brain stimulated in your future studies (my DLPFC is still waiting for theta-burst stimulation)!

Data collection for my main project in the middle of the pandemic would not have been possible without my amazing MSc student Luca. I would like to thank you for your outstanding dedication during these difficult times. I will never forget our exhausting 10-12h shifts in the lab, and the lovely chats we had in between with coffee and cake.

I was lucky to be part of two fantastic research environments. It has been a pleasure to belong to the Neuroplastics group that provided the most welcoming and friendly research environment I could have hoped for. I am also thankful for being part of the Cognitive Computational Neuroscience Journal Club which has inspired me in numerous controversial discussions about the latest research.

Many people have contributed to making my time in Oxford an unforgettable experience. I would like to thank my first housemate, Sam, for controversial debates about statistics and the pandemic, and for eating all my left-over food. My current housemates Momo, Maike, Kumeren and Johny

have enriched my life by creating a welcoming and homely atmosphere in our house share. I hope there will be more opportunities to go on holiday together and celebrate good times. I have found many lovely friends that I enjoy seeing for a walk, coffee or a night out, including Lisa, Patricia and Felicia. Perhaps most importantly, I would like to thank my friends in Berlin, Nubia, Fabienne, Liesa, Thomas, Mandy and Bruno, for being wonderful friends despite the physical distance. Every time I come home and visit you, it feels like I have never left.

Last but for sure not least, I would like to thank the person who convinced me to start my DPhil abroad, my mum. Thank you for supporting me in pursuing my goals throughout my entire life, cooking my favourite meals whenever I come home, and sending enough chocolate to cover any future pandemics.

Table of Contents

List of figures

List of tables

List of abbreviations

AMI	Apathy Motivation Index
ANOVA	Analysis of variance
BDI	Beck Depression Inventory
BIC	Bayes Information Criterion
DLPFC	Dorsolateral prefrontal cortex
DSM-5	Diagnostic and statistical manual of mental disorders
FERT	Facial Expression Recognition Task
fMRI	Functional magnetic resonance imaging
IBLT	Information Bias Learning Task
NIBS	Non-invasive brain stimulation
NNT	Number Needed to Treat
OCI	Obsessive Compulsive Inventory
PET	Positron emission tomography
PILT	Probabilistic Instrumental Learning Task
rTMS	Repetitive transcranial magnetic stimulation
sgACC	Subgenual anterior cingulate cortex
STAI	State-Trait Anxiety Inventory
STICSA	State-Trait Inventory for Cognitive and Somatic Anxiety
tDCS	Transcranial direct current stimulation
TEPS	Temporal Experience of Pleasure Scale
TMS	Transcranial magnetic stimulation

1 Introduction

With a life-time prevalence of around 10-15% depression is one of the most common health conditions worldwide (Lim et al., 2018; Vos et al., 2020). During the Covid-19 pandemic, the rate of depression has increased globally by more than 25% (Santomauro et al., 2021) increasing the need for effective treatments even further. The main symptoms of depression according to the Diagnostic and statistical manual of mental disorders (DSM-5) (American Psychiatric Association & Association, 2013) are depressed mood and loss of interest or pleasure in activities. Depression can also be associated with physical symptoms, such as weight loss or gain, speaking or moving more slowly, and loss of energy, or cognitive symptoms such as feelings of worthlessness or guilt, decreased ability to concentrate or make decisions and suicidal thoughts. The diagnostic criteria for a major depressive episode are met if five of these symptoms, including at least one of the two main symptoms, are present most of the time for at least two weeks and cause significant impairment (American Psychiatric Association & Association, 2013).

The most common treatment approach for depression is antidepressant medication. Antidepressants work by rebalancing neurotransmitter levels in the brain. The remission rate for antidepressants is estimated around 60% (Rush et al., 2006). However, there is large variability in treatment response, and many patients need several attempts before an effective treatment is found. About one third of patients are classified as “treatment-resistant”, i.e. a certain number of treatment attempts have failed (exact number differs between studies)(Rush et al., 2006; Voineskos, Daskalakis, & Blumberger, 2020). Since antidepressants can take 4-6 weeks before showing an effect, finding the right treatment can be a long process. There is therefore a need to develop more effective treatment alternatives.

The development of non-invasive brain stimulation (NIBS) offers a promising new treatment avenue. In modern theoretical accounts of depression, depression is thought to arise from a

dysfunctional interplay of multiple brain networks distributed across the brain. In contrast to other treatment modalities like medications, NIBS can be applied to deliberately target the activity of specific brain regions or networks. NIBS applies magnetic fields or electric currents through the scalp to the brain to increase or decrease excitability in the targeted brain areas which can be used to change the functional interaction of brain networks affected in depression. Transcranial magnetic stimulation (TMS) and transcranial direct current stimulation (tDCS) have been suggested as potential antidepressant treatments. While TMS is approved for treatment-resistant depression in some countries (e.g. U.K., USA) (National Institute for Clinical Excellence, 2015a), tDCS is still under investigation as an antidepressant treatment, with current evidence suggesting small to moderate effect sizes (Razza et al., 2020).

In this thesis, we suggest that the application of NIBS in depression treatment might be improved by including cognitive components – as a measure and/or a manipulation. Cognition plays an important role in depression. Several symptoms of depression are cognitive, such as decreased ability to concentrate and make decisions, or suicidal thoughts. A key concept characterising cognition in depression are negative cognitive biases, i.e. increased processing of negative information. Negative cognitive biases are hypothesised to play a key role in the development and maintenance of depressive symptoms (Beck, 2002). The reduction of negative biases has been shown to precede and predict the clinical response to antidepressant medication, suggesting a causal relationship. Previous research suggests that NIBS can modulate cognition in a way beneficial for depression treatment, for example by reducing vigilance to threat (Ironside et al., 2019; Ironside, O'Shea, Cowen, & Harmer, 2016).

This thesis will address two cognitive approaches to improving NIBS application in depression treatment. First, we propose to apply NIBS in such a way as to target negative cognitive biases directly. More precisely, we suggest that the antidepressant effect of tDCS might be increased by combining it with a learning task designed to reduce negative biases. Second, we propose "cognitive phenotyping" to guide treatment selection for individual patients. There is large

heterogeneity in the symptoms of depression as well as in treatment response, suggesting that different phenotypes with different underlying psychopathology might exist. With respect to TMS treatment, there is emerging evidence that the optimal target region might differ across patients. Previous research has aimed at identifying different phenotypes based on neuroimaging data. As an alternative approach, we propose to define phenotypes based on cognitive task performance. Phenotyping based on each patient's individual cognitive profile might be used to guide treatment selection, such as the choice of TMS target region.

This Introduction chapter is structured into three parts[1]. The first part will review two theories of depression as a network disorder which are important to understand the rationale and mechanisms of NIBS treatment. It will then cover the physiological principles of NIBS and its application in the treatment of depression. Next it will discuss recent research into mechanisms of action of NIBS treatment, and then explain the rationale for these two suggested approaches to improving NIBS treatment. Part 2 will review the evidence for the hypothesis that tDCS might be more effective if combined with a learning task. Part 3 will discuss previous research aimed at personalising NIBS treatment, and the rationale for a phenotyping approach based on cognition.

1.1 Network theories of depression

While early neuroimaging studies investigated the relationship between depression and activity changes in single brain regions, modern network theories of depression state that depression is characterised by a dysfunctional interplay of multiple brain networks distributed across cortical and subcortical areas throughout the entire brain. Two network theories of depression have improved our understanding of, and informed the application of NIBS in depression treatment: The *triple network model of psychopathology* (Menon, 2011), and the *neurobiological model of biased processing of negative stimuli* (Disner, Beevers, Haigh, & Beck, 2011). These models are mutually

[1] Parts of the three introductory chapters have been published in a book chapter (Sarrazin & O'Shea, 2021).

compatible but highlight different aspects of the same underlying theory. While the triple network model emphasises pathophysiology, the biased processing model focuses on dysfunctional information processing. Together, they offer a useful framework for understanding depression as dysfunctional brain network interactions that lead to negatively biased information processing.

According to the *triple network model of psychopathology*, mental health relies on the functional interaction between three different major networks: the executive control network, default mode network and salience network (Figure 1.1). The executive control network is active during task performance and is responsible for executive control processes, such as planning or inhibiting currently irrelevant functions, and for emotion regulation. The executive control network includes the dorsolateral prefrontal cortex (DLPFC), as well as lateral posterior parietal regions. The default mode network is functionally anti-correlated with the executive control network. It is activated at rest and involved in self-referential thinking and episodic memory retrieval. The default mode network spans the medial prefrontal cortex, medial parietal regions, angular gyrus, precuneus, posterior cingulate cortex and posterior hippocampus. The salience network mediates between the executive control and default mode network. It is responsible for the detection of salient stimuli in the environment and enables switching between the other two networks. The salience network includes the dorsal anterior cingulate cortex, amygdala, anterior insula and anterior hippocampus.

In the healthy brain, the salience network reacts appropriately to salient stimuli and activates the executive control network if action in response to the external stimulus is necessary, or activates the default mode network when no interaction with the external world is required. In mental health conditions, the interaction between these three regions is disturbed. Over-activation of the executive control network can result in heightened attention and vigilance to threat, which is characteristic of anxiety disorders. An overactive default mode network, on the other hand, hinders shifting attention away from the internal to the external world. In the framework of the triple network theory, depression is hypothesised to be caused by an overly strong coupling between the

salience network and default mode network, causing excess internal focus and consequent symptoms such as loss of interest in activities and rumination.

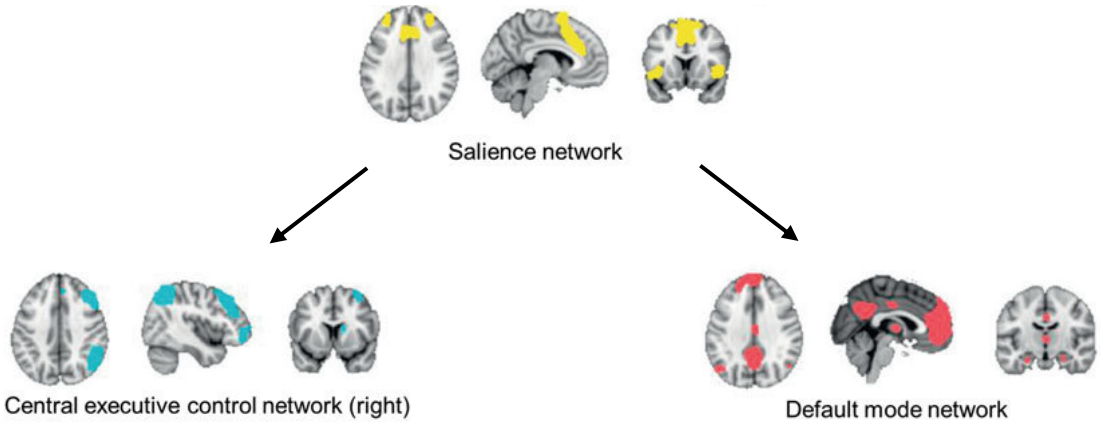

Figure 1.1. The three major brain networks critical to mental health according to the triple network model of psychopathology. The model states that the salience network reacts to salient stimuli in the environment controlling the balance between the default mode network activated at rest, and the executive control network activated during interaction with external stimuli. Depression is supposed to origin from increased functional correlation between the salience network and default mode network which results in a shift from interaction with the outside world towards increased attention to internal processes which can lead to depressive symptoms such as rumination. Reproduced from (Menon, 2011).

The *neurobiological model of biased processing of negative stimuli* is based on Beck's cognitive model of depression (Disner et al., 2011)(Figure 1.2). This model outlines how dysfunctional interactions between the executive control network (DLPFC), default mode network (which the subgenual anterior cingulate cortex is connected to) and salience network (dorsal anterior cingulate cortex, amygdala) are hypothesised to contribute to the development and maintenance of depressive symptoms. According to this model, negative stimuli elicit a heightened response in the thalamus. This leads to pathologically high activity in the amygdala and the subgenual cingulate cortex. Cortical control areas such as the prefrontal cortex are hypo-active in depression so that they fail to down-regulate the response of the limbic areas (including amygdala, thalamus and subgenual cingulate which form a network involved in emotional responses) to negative stimuli. Overall, this dysfunctional circuit leads to increased attention to and processing of negative stimuli.

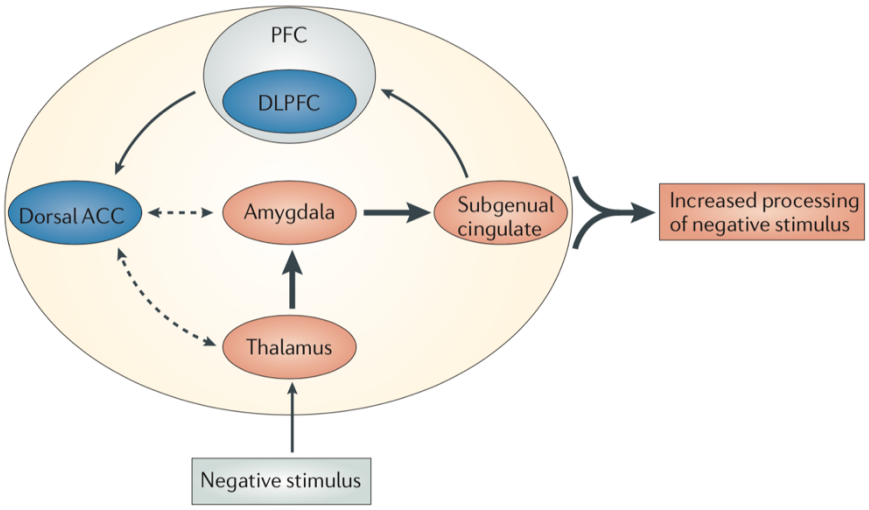

Figure 1.2. Cognitive neurobiological model of biased processing of negative stimuli. In depression, increased processing of negative stimuli is hypothesized to arise from increased activity in the thalamus, amygdala and subgenual cingulate cortex in response to negative stimuli, accompanied by decreased regulatory influence from the dorsal anterior cingulate cortex and dorsolateral prefrontal cortex (DLPFC) on these interconnected subcortical regions. Hyperactive areas are shown in red, hypoactive areas in blue. Dashed arrow indicate attenuated connectivity, thicker arrows indicate increased connectivity. Reproduced from (Disner et al., 2011).

Both models identify the same two core pathological features in depression: First, hyper-active bottom-up emotional drive from the limbic system (and subgenual anterior cingulate cortex) which is reflected in hyper-connectivity between the salience network and default mode network; and second, hypometabolism in DLPFC which leads to reduced top-down cortical regulatory influence on limbic activity and is associated with hypo-activity of the executive control network.

Although the first approach of NIBS depression treatment was developed before the view of depression as a network disorder became predominant (see 1.2.1.2 and 1.2.2.2), these two network theories of depression frame the current conceptual approach to NIBS treatment in depression: 1) Depression does not arise from dysfunction of discrete brain regions, but from dysfunctional interplay within and between distributed cortico-subcortical brain networks. This perspective suggests that the DLPFC, which is usually stimulated in depression treatment, is only one out of several potential target regions; 2) A key feature of depression pathophysiology is abnormal functional interactions between cortical and subcortical areas. More precisely, subcortical limbic

regions seem to overreact in response to negative stimuli, while prefrontal cortical areas show impaired downregulation of those negative emotional responses (Disner et al., 2011). Hence, restoring normal functional interactions within cortico-limbic circuits may be an important mechanism of effective treatment action; 3) Within the triple network perspective, hyperconnectivity between subgenual ACC and the default mode network, which is associated with rumination, seems to be a key psychopathological feature of depression (Greicius et al., 2007).

1.2 Non-invasive brain stimulation as novel treatment approach for depression

NIBS applies electric currents or magnetic fields to change the excitability in target regions which in turn leads to physiological and functional changes. NIBS is a promising approach for depression treatment since it can be used to up- or downregulate activity in the networks affected in depression. In contrast to antidepressant medication, specific (cortical) brain regions can be targeted. NIBS treatment has been shown to be effective for a proportion of patients who do not respond to conventional first-line treatments. Transcranial magnetic stimulation (TMS) applied to the left DLPFC is an approved antidepressant treatment in some countries, including the UK and the US. Transcranial direct current stimulation (tDCS), a safer, cheaper and more accessible alternative, is currently under investigation as antidepressant treatment. The following sections will discuss the physiological bases and clinical application of TMS (1.2.1) and tDCS (1.2.2) in depression.

1.2.1 Transcranial magnetic stimulation

1.2.1.1 Physiological effects of TMS

TMS uses focal magnetic pulses generated by rapidly changing current flowing through the stimulating coil. The magnetic fields pass through the scalp and induce an electric current in the

brain which is strong enough to depolarise neurons and induce action potentials in cortical layers. The physiological effects depend on the type and orientation of the coil, the intensity and frequency of the stimulation, the waveform of the magnetic pulses, as well as the relative orientation of the neurons (Kammer, Beck, Thielscher, Laubis-Herrmann, & Topka, 2001; Lefaucheur et al., 2014; Sommer et al., 2006). The most commonly used coil type is the "figure-of-eight" shaped coil which depolarises brain tissue in a volume of a few cubic centimetres (Figure 1.3). The highest magnetic field strength is achieved in the middle of the coil where the windings meet which leads to high focality. The direct physiological effects of TMS are mostly restricted to superficial cortical layers. However, TMS can reach distal brain regions through indirect effects via connecting fibres from the stimulated cortical region. To increase the effect on subcortical targets, novel coil shapes have been created that induce currents in deeper brain layers while still maintaining focality.

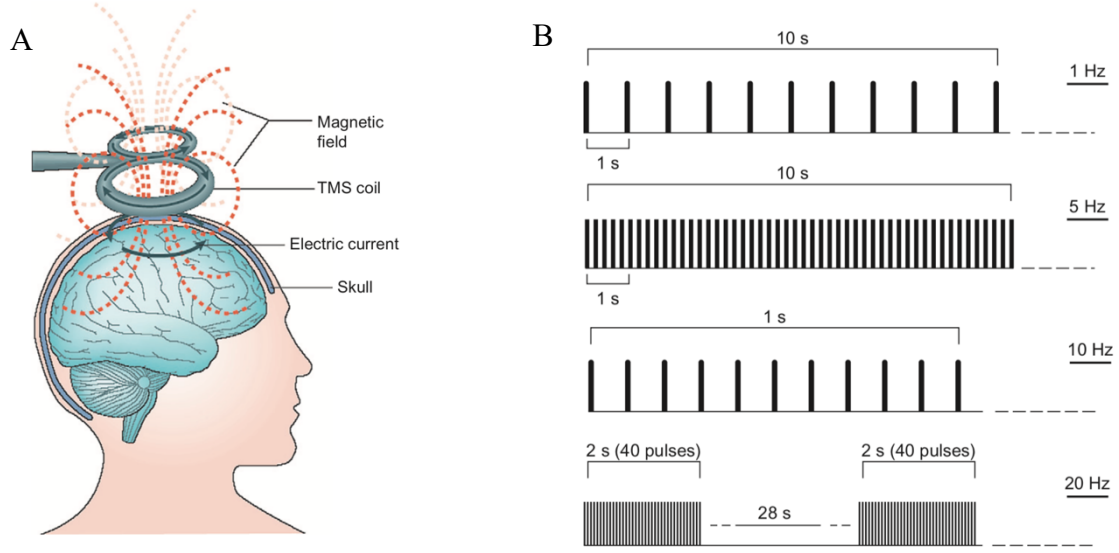

Figure 1.3. Set-up and protocols for TMS. (A) TMS applies magnetic fields transcranially to induce electric currents in the brain. (B) Common rTMS protocols: 10s of low frequency rTMS at 1 Hz (first row) or high-frequency rTMS at 5Hz (second row), 1s of high-frequency TMS at 10 Hz (third row), and high-frequency TMS at 20Hz applied in trains of 2s (fourth row). Figures retrieved from (Spronk, Arns, & Fitzgerald, 2011).

TMS can be applied in single pulses or in several consecutive pulses which is referred to as repetitive TMS (rTMS)(Figure 1.3). RTMS is of clinical interest since it has long-lasting effects on

cortical excitability – effects outlast the stimulation period. Physiological effects of TMS have mainly been investigated in the motor cortex since it provides a direct measurable readout. Cortical excitability can be investigated by stimulating a specific subregion of the motor cortex corresponding to the hand representation within the motor cortex which induces muscle activity in the contra-lateral hand muscle. Muscle activity is measurable as waveforms, so-called motor-evoked potentials, that can be recorded via electrodes placed on the hand muscle. Changes in the amplitude of motor evoked potentials are a measure for cortical excitability, e.g. an increase in cortical excitability will lead to motor evoked potential with larger amplitudes. High-frequency rTMS (≥ 5Hz) has been shown to increase, whereas low-frequency rTMS (≤ 1Hz) has been shown to decrease the amplitude of motor evoked potentials (Chen et al., 1997; Pascual-Leone, Valls-Solé, Wassermann, & Hallett, 1994).

These frequency-dependent effects on cortical excitability are assumed to generalise to other cortical regions. The physiological effects on other brain areas are more difficult to investigate due to the lack of direct physiological readout such as motor evoked potentials. The first evidence for frequency-dependent TMS effects in the DLPFC came from positron emission tomography (PET) studies which investigated how brain metabolism changes in response to one or multiple TMS sessions. These studies showed that high-frequency TMS increases, whereas low-frequency TMS decreases metabolism in the DLPFC, the standard target region in NIBS depression treatment (Loo et al., 2003; Speer et al., 2000).

Importantly, the effects of rTMS on cortical excitability can outlast the stimulation period by minutes up to hours. For example, a single session of high-frequency rTMS can increase the amplitude of motor evoked potentials for up to one hour post-TMS offset (Ziemann et al., 2008). The long-lasting effects are hypothesised to rely on plasticity mechanisms similar to long-term potentiation (high-frequency rTMS) and long-term depression (slow-frequency rTMS)(Cooke & Bliss, 2006; Huang, Lu, et al., 2017). To induce antidepressant effects, TMS is applied in repeated

sessions so that the effects accumulate over time, presumably via long-lasting changes in the physiological and functional interaction of brain networks affected in depression.

1.2.1.2 TMS for the treatment of depression

The first studies that applied TMS as an intervention in depression were motivated by findings from PET studies which showed that depression was associated with decreased metabolism in the left DLPFC (George, Ketter, & Post, 1994). This led to the hypothesis that excitatory high-frequency TMS applied to the DLPFC could potentially increase local cortical excitability in the target region and thus possibly result in reduced depressive symptoms. In the first blinded and controlled study, 10 Hz TMS applied to the left DLPFC in 5 daily sessions had antidepressant effects compared to sham TMS and TMS applied to a control region (Pascual-Leone, Rubio, Pallardó, & Catalá, 1996). Subsequently, multiple clinical trials were conducted (George et al., 2010; Levkovitz et al., 2015; O'Reardon et al., 2007) which ultimately led to the approval of TMS as a treatment for treatment-resistant depression by the U.S. Food and Drug Administration in 2008, and later by the National Institute for Clinical Excellence in the UK in 2015 (National Institute for Clinical Excellence, 2015a). The recommended and most commonly used protocol is the application of 3,000 pulses per session at a frequency of 10Hz in trains of 4 seconds at an intensity of 120% of resting motor threshold to the left DLPFC (Perera et al., 2016). The treatment duration is 4-6 weeks with five sessions per week.

As an alternative to high-frequency stimulation of the left DLPFC, low-frequency TMS applied to the right DLPFC has also been investigated (Berlim, Van den Eynde, & Jeff Daskalakis, 2013). This is based on a theory that depression might result from a dysfunctional asymmetry in cortical activity between left and right hemisphere, especially left and right DLPFC (Grimm et al., 2008; Hecht, 2010). According to this theory, the left hemisphere, which is primarily associated with processing of positive information, is thought to be hypoactive, whereas the right hemisphere, associated with the processing of negative information, is thought to be hyperactive (see Hecht

(2010) for a review). Consistent with this framework, although high-frequency TMS applied to the left DLPFC is the standard treatment, a large meta-analysis showed that low-frequency TMS applied to the right hemisphere is equally effective (Lepping et al., 2014).

The response rate for high-frequency TMS applied to the left DLPFC is around 47% (Number Needed to Treat (NNT) = 6). The remission rate is around 27% (NNT = 8) (Berlim, van den Eynde, Tovar-Perdomo, & Daskalakis, 2014; Blumberger et al., 2018). Importantly, the patients recruited in TMS depression trials are usually relatively severe cases of depression classified as "treatment-resistant", i.e. patients had previously undergone multiple unsuccessful treatment attempts.

1.2.2 Transcranial direct current stimulation

TDCS has been proposed as a safer, cheaper and more accessible alternative to TMS for the treatment of depression. Since tDCS does not induce action potentials in the brain, there is a lower risk of seizures compared to TMS. Moreover, tDCS is a portable device, so would be suitable for home use (Alonzo et al., 2019; Shaw et al., 2017). This makes tDCS accessible to a much larger number of patients compared to TMS which is a large, heavy, expensive and non-portable device and needs to be administered by trained clinical or research staff.

1.2.2.1 Physiological effects of tDCS

In contrast to TMS, tDCS applies diffuse electric currents to the brain. The current is delivered via electrodes that are fixed on the scalp. Animal research has shown that tDCS acts by shifting the membrane potential of cortical neurons (Bikson et al., 2004; Bindman, Lippold, & Redfearn, 1962; Rahman et al., 2013). The direction of the effect (i.e. depolaristion vs. hyperpolaristion) depends on the direction of the current flow, with respect to the orientation of the neurons (Figure 1.4A). Anodal stimulation (inflowing current) applied in parallel to the axis of pyramidal neurons

depolarises the soma facilitating the occurrence of action potentials. In contrast, cathodal tDCS (outflowing current) hyper-polarises the soma making it less likely for action potentials to occur. The net effect that tDCS has on neurons thus depends on the polarity of the current (anodal vs. cathodal) and the orientation of the neurons (Liu et al., 2018).

Since the cortex in the human brain is folded, tDCS produces a complex pattern of anodal and cathodal stimulation (Karabanov, Saturnino, Thielscher, & Siebner, 2019). The current typically enters a gyrus on one side (anodal stimulation), and leaves it on the other side (cathodal stimulation), so that the same gyrus receives anodal as well as cathodal stimulation (Figure 1.4B). The proximity of a brain region to the anode and cathode, respectively, determines whether anodal or cathodal stimulation predominates. Since the effects of tDCS are dependent on brain anatomy, there is large inter-individual variability in the direct physiological effects (Opitz, Paulus, Will, Antunes, & Thielscher, 2015) as well as in the resulting excitability changes (Wiethoff, Hamada, & Rothwell, 2014).

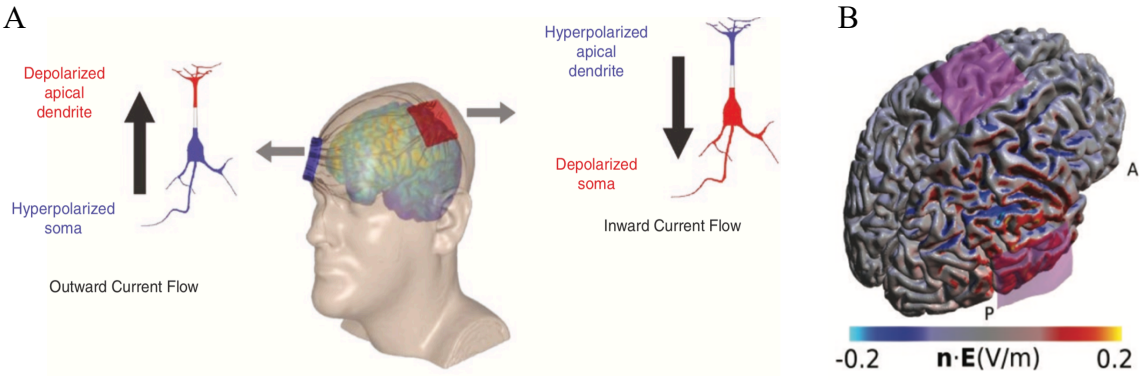

Figure 1.4. Physiological effects of tDCS. (A) Physiological effects of anodal (red) and cathodal (blue) tDCS on the soma and apical dendrite. Reproduced from (Sharma, Farahani, Bikson, & Parra, 2021). (B) Example of the complexity of electric fields induced by tDCS. Since the electric current tends to enter a gyrus on one side and leave it on the other, tDCS induces a complex pattern of anodal (red) and cathodal (blue) stimulation. Reproduced from (Karabanov et al., 2019).

Similar to TMS, the physiological effects of tDCS in humans have mainly been investigated in the motor cortex since motor-evoked potentials provide a convenient indirect measure of changes in

cortical excitability. Anodal tDCS has been shown to increase, whereas cathodal tDCS has been shown to decrease the amplitude of motor-evoked potentials, effects which out-last the stimulation period (Bergmann et al., 2009; Nitsche & Paulus, 2000). Neuroimaging studies have confirmed that tDCS applied to the DLPFC also induces polarity-specific effects in the DLPFC, e.g. increased perfusion in response to anodal tDCS (Stagg et al., 2013).

Similar to rTMS, the polarity-dependent effects of tDCS on cortical excitability can outlast the period of stimulation by minutes up to hours (Bergmann et al., 2009; Nitsche & Paulus, 2000). These long-lasting effects are hypothesised to rely on mechanisms of synaptic plasticity. Direct evidence comes from a study in which direct current stimulation applied to a mouse brain slice induced long-lasting synaptic potentiation similar to long-term potentiation, in a polarity-specific way (Fritsch et al., 2010). Consistent with this, anodal tDCS applied to the motor cortex facilitated motor skill acquisition in humans. Because of these plasticity induction effects, tDCS might have the potential to restore the normal functional interaction of the brain networks affected in depression.

1.2.2.2 TDCS for the treatment of depression

Neurophysiological effects of current applied to the scalp were already discovered in the 18th century (Brunoni et al., 2012) but in early studies tDCS did not show reliable effects in the treatment of depression in humans (Arfai, Theano, Montagu, & Robin, 1970). Around the time that the first successful TMS trials were conducted in the late 20th century, tDCS was rediscovered. The first studies in human participants showed that tDCS has effects on motor cortico-spinal excitability depending on the polarity (Nitsche & Paulus, 2000; Priori, Berardelli, Rona, Accornero, & Manfredi, 1998). The application of tDCS in depression treatment built on the same rationale as TMS treatment, i.e. it aimed to increase excitability of the left DLPFC (Fregni et al., 2006). However, evidence to date for the efficacy of tDCS as an antidepressant treatment is mixed. While some clinical trials have reported antidepressant effects comparable to other treatments

(Brunoni et al., 2016; Brunoni et al., 2013) tDCS has been inferior to antidepressant drug outcomes in other studies (Brunoni et al., 2017).

Therefore, more evidence of efficacy is needed before tDCS might be approved as an antidepressant treatment (National Institute for Clinical Excellence, 2015b). Several clinical trials using tDCS on its own or in conjunction with other treatment modalities are currently being conducted. The response rate for tDCS treatment in depression is currently estimated at 34% (NNT = 7), with a remission rate of 23.1% (NNT = 9)(Brunoni et al., 2016). When comparing efficacy across TMS and tDCS trials, it should be noted that while TMS trials usually focus on patients with treatment-resistant depression, tDCS trials often recruit a wider range of patients including less severe cases.

1.3 Mechanisms of action of non-invasive brain stimulation treatment

As outlined in the previous sections, recent theories of depression posit that depression arises from a disturbed interplay of the three major brain networks. NIBS has the potential to modulate the functional interaction of brain networks by increasing or decreasing the excitability of cortical regions and anatomically connected subcortical regions. In the past decade, important insights into the mechanisms of action of NIBS treatment have been gained in neuroimaging studies. In most neuroimaging studies, participants were scanned before and after undergoing NIBS treatment, and changes in brain function before vs. after treatment were analysed. Potential candidate mechanisms of NIBS treatment action could thus be identified by correlating changes in brain function to clinical improvement.

The vast majority of studies have focused on resting-state functional connectivity which is used to investigate the functional interaction of brain networks in the absence of any task. Connectivity is especially relevant to TMS since the magnetic field induced by TMS is focal, and changes in remote regions mainly occur indirectly through anatomical connections from the stimulated brain

region. Resting-state functional connectivity analyses can be used to investigate how focal stimulation of a cortical target is propagated within and across interconnected networks. Resting-state functional magnetic resonance imaging (fMRI) has some advantages over task-related fMRI, including better signal-to-noise ratio, no confounds from task-related aspects, short acquisition time and higher participant compliance (Fox & Greicius, 2010). Resting-state connectivity measures may be particularly important in treatment studies of depression. In a healthy brain, the default mode network, normally deactivated during task states, becomes active at rest. In depression, it is hypothesised to be hyperactive at rest and this is associated with rumination (Berman et al., 2011; Greicius et al., 2007). Dampening down default mode network hyperactivity is a candidate mechanism underlying NIBS treatment, similar to other treatment modalities that have been shown to have this effect (eg: pharmacological (Brakowski et al., 2017)).

1.3.1 Mechanisms of action of TMS treatment

In line with the *neurobiological model of biased processing of negative stimuli,* several studies suggest that TMS not only increases activity in the stimulated DLPFC, but also affects resting-state functional connectivity of the DLPFC to subcortical regions involved in emotion regulation, such as the amygdala (Eshel et al., 2020; Richieri et al., 2017; Salomons et al., 2014). Many studies have focused on connectivity to the subgenual anterior cingulate cortex (sgACC) based on findings highlighting the role of the sgACC in depression (Greicius et al., 2007; Mayberg et al., 2005). Fox and colleagues hypothesised that differences in the clinical efficacy of TMS across individuals might be related to differences in the connectivity strength between each individual's target region and the sgACC (Fox, Buckner, White, Greicius, & Pascual-Leone, 2012)(Figure 1.5). They retrospectively analysed data from several previous TMS trials and compared connectivity of the target region with clinical efficacy across studies. Resting-state connectivity of each target region was estimated based on group-averaged connectivity in a large sample of healthy individuals. They found that across individuals and studies stronger treatment response was associated with DLPFC target regions that displayed a stronger negative correlation to the sgACC (Figure 1.5). In line with

the triple network theory of psychopathology, this study could be interpreted to suggest that the antidepressant effect of TMS might be related to downregulation of activity within the sgACC and/or default mode network. Several subsequent studies have reported changes in sgACC connectivity in response to TMS treatment (Baeken, Duprat, Wu, De Raedt, & van Heeringen, 2017; Baeken et al., 2014; Liston et al., 2014; Philip et al., 2018; Salomons et al., 2014; Taylor et al., 2018).

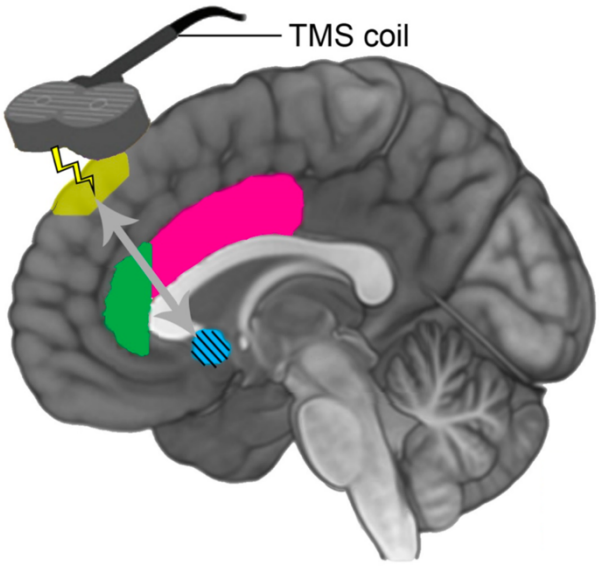

Figure 1.5. Effect of TMS on subcortical regions. TMS has been shown to be more effective the stronger the negative correlation between the stimulated region and the subgenual anterior cingulate cortex (Fox et al., 2012). Yellow: DLPFC, blue: subgenual anterior cingulate cortex, pink: dorsal ACC, green: rostral ACC. Figure adapted from (Pantazatos et al., 2022).

Based on findings that depression is associated with disturbed interactions between brain networks, a few studies have investigated how TMS treatment changes connectivity within and between the three key networks implicated in depression - default mode network, executive control network and salience network. In one study, the signal within two seed regions, the DLPFC and sgACC, was correlated with regions in the default mode network and executive control network (Liston et al., 2014). TMS reduced hyperconnectivity of the sgACC to the default mode network and induced negative connectivity between DLPFC and medial prefrontal cortex. In another study, positive clinical outcome was associated with decreased connectivity between sgACC and default mode

network regions, as well as reduced connectivity between the hippocampus and regions of the salience network (Philip et al., 2018).

To sum up, the modulation of subcortical limbic areas related to emotional response seems to play an important role in the antidepressant effect of NIBS treatment. Therefore, connectivity between the cortical target area and subcortical regions (in particular the sgACC) is especially relevant. Resting-state network analysis suggests TMS treatment acts by modulating the functional interaction of the three major resting-state networks, especially by downregulating connectivity in the default mode network.

1.3.2 Mechanisms of action of tDCS treatment

Research on the mechanisms of tDCS treatment is in the early stages, and most studies were conducted in healthy individuals. There is evidence that tDCS modulates the functional interaction of the three major brain networks as postulated by the triple network theory of psychopathology. In two studies, participants underwent resting-state fMRI directly after the application of a single tDCS session. In the first study, real compared to sham tDCS induced changes in connectivity close to the electrodes as well as in more remote regions - the default mode network and bilateral frontoparietal network (Keeser et al., 2011). In the second study, active tDCS led to increased synchrony in the executive control network and decreased synchrony in the default mode network (Pena-Gomez et al., 2012). Another study investigated the effects of prefrontal tDCS on resting-state networks using arterial-spin labelling (Stagg et al., 2013). During anodal stimulation, tDCS induced increases in perfusion in areas close to the stimulation site. After the stimulation, a decrease in regions associated with the default mode network was observed.

A recent modelling study suggests that tDCS might actually have a direct effect on the default mode network (Karabanov et al., 2019). The authors found that the highest electric field strength is generated in between the electrodes, i.e. in the medial prefrontal cortex rather than the dorsolateral

prefrontal cortex. Given the important role of sgACC and default mode network connectivity in NIBS treatment, this might suggest that the tDCS setup chosen to stimulate the DLPFC might actually act by stimulating the default mode network directly.

To conclude, the limited number of studies conducted to date demonstrate that DLPFC-targeted tDCS decreases activity in the default mode network and increases activity in the executive control network, consistent with the triple network theory of antidepressant effects. However, these studies were conducted in healthy participants and measured the effects of a single tDCS session. Hence, it remains unclear both what the long-term effects of a daily tDCS treatment protocol over several weeks would be, and how those physiological effects would relate to clinical improvement.

Overall, there is convincing evidence that TMS as well as tDCS modulate the functional interaction between the three major resting state networks. Specifically, NIBS seems to upregulate activity in the executive control network, and downregulate activity in the default mode network. This suggests several alternative candidate regions as potential treatment targets beyond the conventional DLPFC. Cortical areas closely coupled with the default mode network, and that are easy to reach with NIBS (eg: medial prefrontal cortex or angular gyrus) might be good alternative targets for depression treatment. Two recent studies created maps of cortical areas which are most strongly connected to the networks affected in depression and constitute potential alternative target regions (Siddiqi et al., 2021; Zhang et al., 2020)(Figure 1.6).

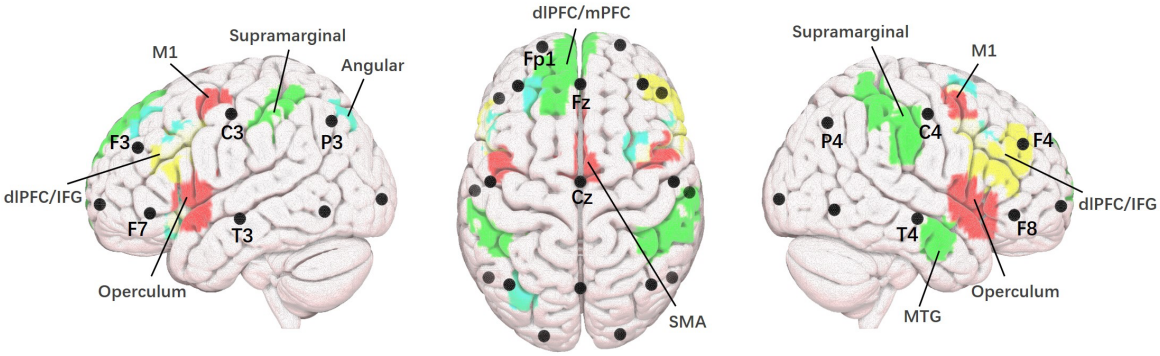

Figure 1.6. Map of potential cortical target regions connected to areas associated with depression. The different colours refer to three different analysis pipelines. Yellow: cortical regions that are part of the clusters associated with depression. Light blue: cortical areas most strongly correlated to the regions associated with depression (positive correlation). Red and green: Cortical regions correlated with the highest number of depression-associated areas (red: positive correlation, green: negative correlation). Figure provided by Binlong Zhang, adapted from (Zhang et al., 2020).

1.4 Improving the application of NIBS in depression treatment

In the past decade, a lot of progress has been made in the understanding of mechanisms of NIBS treatment on the neurophysiological level. A lot of research has aimed to identify the key neurophysiological changes related to clinical improvement, and determine the optimal target region (mostly within DLPFC). However, very little research has been conducted on the cognitive level. Cognition is relevant to improving the application of NIBS in depression treatment for two reasons: 1) Depression is characterised by multiple cognitive symptoms, and negatively biased information processing is hypothesised to be a key causal factor in the development and maintenance of depressive symptoms. 2) There is evidence that NIBS can modulate cognitive effects that might be beneficial in the treatment of depression, for example by reducing vigilance to threat (Ironside et al., 2016). Therefore, this thesis will address two potential cognitive approaches to improving the application of NIBS in the treatment of depression.

First, we suggest that the antidepressant treatment efficacy of tDCS might be increased by combining it with a learning task targeting negative biases. In clinical trials, tDCS is usually

applied at rest. However, tDCS can modulate synaptic plasticity and learning effects. The effect of tDCS might therefore be increased by applying it during a learning task that is relevant to depression, rather than by applying it during rest. Depression is associated with negative cognitive biases which are hypothesised to reflect underlying deficits in learning the statistics of rewards and punishments. We therefore chose to combine tDCS with a reinforcement learning task. The goal of this project was to establish proof of concept that tDCS applied during reinforcement learning would have a greater functional impact than tDCS applied during rest.

Second, to address the problem of variability in NIBS treatment response, we reasoned that patient stratification based on different cognitive phenotypes of depression could prove useful in guiding treatment selection. Depression is associated with large interindividual variability – regarding symptomatology as well as treatment response. The variability in symptoms and treatment response suggests that different phenotypes with different underlying psychopathology might exist. Therefore, much effort has attempted to define different subtypes of depression, grouping together individuals with similar psychopathology who might respond to the same treatment. While most research has investigated phenotypes based on clinical symptoms or neuroimaging markers, we propose that phenotypes could be defined based on cognitive task performance. Cognitive phenotyping could be used to guide treatment selection, including (but not limited to) the choice of target area for TMS treatment.

2 Targeting negative cognitive biases with tDCS

In the clinical trials conducted to date, tDCS has been applied at rest. However, tDCS might be more effective if applied during a learning task, since tDCS has been shown to induce plasticity and learning effects (Fritsch et al., 2010; Reis et al., 2009). In order to improve its efficiency, tDCS could be combined with a learning task designed to reduce negative cognitive biases. Negative biases are theorised to play a major role in the maintenance of depressive symptoms (Kube, Schwarting, Rozenkrantz, Glombiewski, & Rief, 2020) and reversal of negative biases is thought to be one mechanism underlying effective treatments (Disner et al., 2011). Recent work indicates that tDCS might have the potential to reduce negative biases (Clarke, Browning, Hammond, Notebaert, & MacLeod, 2014; Ironside et al., 2019; Ironside et al., 2016). This suggests the hypothesis that antidepressant-like effects of tDCS might be increased if applied during a learning task designed to reduce negative biases.

In this thesis two proof-of-concept studies in healthy individuals (Chapter 4) and individuals with low mood (Chapter 6) were conducted to test whether tDCS can reduce negative biases. The first part of this section will discuss how negative biases in depression might develop from a computational perspective. The second section will discuss evidence supporting the hypothesis that tDCS might have the potential to reduce negative biases.

2.1 Negative biases in depression

Depression is characterised by a negative cognitive bias, i.e. cognition and attention are biased towards negative rather than positive information. Compared to healthy controls, individuals with depressive symptoms remember more negative words (Bradley, Mogg, & Williams, 1995), perceive feedback as more negative (Gotlib, 1983), tend to interpret ambiguous information as

negative (Everaert, Podina, & Koster, 2017) and overestimate their likelihood of experiencing adverse life events (Korn, Sharot, Walter, Heekeren, & Dolan, 2014).

Negative biases are theorised to play a key role in the development and maintenance of depressive symptoms. According to Beck's cognitive model of depression (Beck, 1967) stressful experiences activate depressive schemas which are characterised by negative self-referential beliefs. These depressive schemas bias attention, information processing and memory towards negative information, which in turn reinforces the depressive schema. Once initiated, this feedback loop leads to the development and maintenance of depressive symptoms. In this model, negative biases are thus hypothesised to play a causal role in depression (Disner et al., 2011). Empirical evidence for a causal role of negative biases in depression comes from a study on antidepressant medication (Godlewska, Browning, Norbury, Cowen, & Harmer, 2016). A reduction in negative biases was already observed after one week of treatment with an antidepressant, i.e. before clinical symptoms improved. Importantly, the reduction of negative biases after one week of treatment predicted who went on to respond to the treatment. These findings suggest that negative biases play a causal role in depression, and a reduction in negative biases might be a mechanism underlying antidepressant treatments.

2.1.1 Computational modelling

Important insights into how affective biases develop have been gained in the field of computational psychiatry. One aim of computational psychiatry is to identify the mechanisms in information processing that lead to the development of mental disorders. Using computational modelling, it is possible to distinguish different latent processes that are involved in information processing and decision making. Affective biases can be captured in reinforcement learning models which describe how agents process rewards and/or punishments to optimise future choices. For example, participants may be presented with a number of options which differ in their probability of being rewarded (or punished). Often, the probability structure of the task is not explicitly explained to the

participants so that they need to infer the best option by trial-and-error over the course of many trials.

Reinforcement learning models commonly include a Rescorla-Wagner learning rule (Rescorla & Wagner, 1972). The agent is assumed to learn the expected value of each option based on the outcome history. The value of option *k*, Q^k, is updated after each trial *t* following this rule:

$$Q_{t+1}^k = Q_t^k + \alpha(r_t - Q_t^k)$$

The value of option k is updated with a proportion of the prediction error, i.e. the difference between the expected outcome Q_t^k and the observed outcome r_t. The speed of updating is determined by the learning rate parameter α. A low learning rate means that the agent integrates outcomes over several timepoints and changes their estimated value only slowly when the underlying reward association changes. A higher learning rate reflect a larger influence of the most recent outcomes so that behaviour changes more quickly in response to changes in the reward structure.

In a reinforcement learning framework, agents are assumed to keep track of the value of each option and use these estimates to optimise their choices. To transform value estimates into action probabilities, a Softmax function is commonly applied. The probability of the agent to choose option *k* on trial *t*, p_t^k, corresponds to:

$$p_t^k = \frac{\exp(\beta Q_t^k)}{\sum_{i=1}^{K} \exp(\beta Q_t^i)}$$

The inverse temperature parameter β determines the slope of the Softmax function and captures choice stochasticity, i.e. how likely participants are to choose the option with the highest value. A high inverse temperature means that participants are likely to choose the better option even if the value difference between the options is small. A low inverse temperature indicates that participants

are less likely to choose the best option, either because they have difficulties keeping track of the reward probabilities, or because they use a different strategy which is not captured by the model.

Some models include an additional parameter in equation (1), the reward sensitivity parameter *p* which scales the outcome r_t. This parameter is of particular interest in depression research since it can be used to capture differences in sensitivity to rewards vs. punishments. However, it can be challenging to distinguish the reward sensitivity parameter from the inverse temperature (Huys, Pizzagalli, Bogdan, & Dayan, 2013)

2.1.2 Development of negative biases in a computational framework

Research on reinforcement learning in depression has aimed at identifying alterations in information processing associated with depression that might contribute to negative biases. Although associations with depression have been observed for several different computational parameters, two hypotheses have predominated: (1) Negative biases might arise from an imbalance in sensitivity to rewards vs. punishments (Huys et al., 2013; Pike & Robinson, 2022), and (2) negative biases might arise from an imbalance in reward vs. punishment learning rates (Aylward et al., 2019; Eshel & Roiser, 2010; Pike & Robinson, 2022). While both mechanisms might contribute to negative biases in depression, they suggest different treatment approaches. An imbalance in reward vs. punishment sensitivity could be treated by targeting how positive (or negative) events are perceived. If negative biases are caused by increased punishment learning rates, individuals could instead be taught not to change their behaviour immediately in response to negative feedback. Therefore, understanding which computational mechanism leads to negative biases can help to optimise treatment approaches.

While some research indicates that depression might be characterised by an increased sensitivity to punishments and/or decreased sensitivity to rewards (Huys et al., 2013; Pike & Robinson, 2022; Pizzagalli et al., 2009; Pizzagalli, Jahn, & O'Shea, 2005; Taylor Tavares et al., 2008), there is also

contradicting evidence (Aylward et al., 2019). More recent research has focused on alterations in learning rates, i.e. the way individuals learn from unexpected information and update their beliefs accordingly. Depression has been linked to increased punishment learning rates (Aylward et al., 2019; Beevers et al., 2013; Huang, Thompson, & Paulus, 2017; Pike & Robinson, 2022) and decreased reward learning rates (Brown et al., 2021). This imbalance in punishment vs. reward learning suggests that individuals suffering from depression change their behaviour more quickly in response to unexpected negative outcomes compared to unexpected positive events than healthy individuals. In real life, this might cause individuals to feel only little encouragement by positive feedback and give up very quickly after receiving negative feedback. This might lead to a loop of avoiding situations once they have been associated with one negative outcome so that the negative experience cannot be outweighed by future positive experiences.

While initial studies suggested that a negative bias might be associated with relatively higher punishment than reward learning rates, recent research suggests that a key mechanism underlying affective biases might lie in adjusting learning rates appropriately to the context (Browning, Behrens, Jocham, O'Reilly, & Bishop, 2015; Gagne, Zika, Dayan, & Bishop, 2020). One feature that determines whether a higher or lower learning rate is more appropriate is the instability, or volatility, of the underlying reward association (Behrens, Woolrich, Walton, & Rushworth, 2007). In a volatile environment, the strength of the association between actions and outcomes changes over time. Unexpected outcomes in a volatile environment are more likely to reflect an actual change in reward probabilities than in a stable environment, i.e. are more informative for optimising choice behaviour. The challenge for the agent is to determine whether an unexpected outcome was caused by a change in the underlying reward associations, or by noise (i.e. because the association between actions and outcomes is not 100%). In the former case, the agent should change their behaviour to optimise future choices, whereas in the latter case the agent should stick with their previous strategy.

Human participants have been shown to adjust their learning rates to the volatility of the environment in the expected way, i.e. they show higher learning rates in volatile than in stable conditions. Anxiety and depression have been associated with deficits in adjusting learning rates to the volatility of the context (Browning et al., 2015; Gagne et al., 2020). Some evidence suggests that this might be especially the case for aversive learning (Browning et al., 2015). Since learning rates reflect the perceived informativeness of events, deficits in adjusting learning rates to the context suggest that depression might be associated with a misestimation of informativeness. A negative bias might naturally arise if individuals estimate negative events to be more informative than positive events.

A recent study tested whether the human brain is able to maintain separate estimates of the informativeness of positive and negative events, which could provide the basis for an affective bias (Pulcu & Browning, 2017). The authors developed a new paradigm which manipulated the volatility (i.e. informativeness) of wins and losses separately (Figure 2.1A). The data were analysed using a computational model with separate learning rates for win (positive) and loss (negative) outcomes. Participants adjusted their learning rates to the volatility of the respective outcome, i.e. when wins were volatile, win learning rates were higher than loss learning rates, and vice versa when losses were volatile (Figure 2.1B). These findings confirmed the hypothesis that humans are able to maintain separate estimates of the informativeness of positive and negative information and adjust their behaviour accordingly, i.e. humans show affective biases towards the type of information (positive or negative) which they estimate to be more informative. Given that depression has been associated with difficulties in adjusting behaviour to changes in informativeness, a negative bias might arise from overestimation of the informativeness of negative events.

To summarise, the findings described above suggest that negative biases in depression might be caused by deficits in reinforcement learning. More precisely, a current hypothesis states that individuals suffering from depression might overestimate how useful negative events are in

optimising future choices, which might in turn increase the focus on negative compared to positive information. This theory suggests a potential intervention: if negative biases result from overestimation of informativeness of negative events, a training paradigm in which positive events are more informative than negative events might be able to counter-act negative biases.

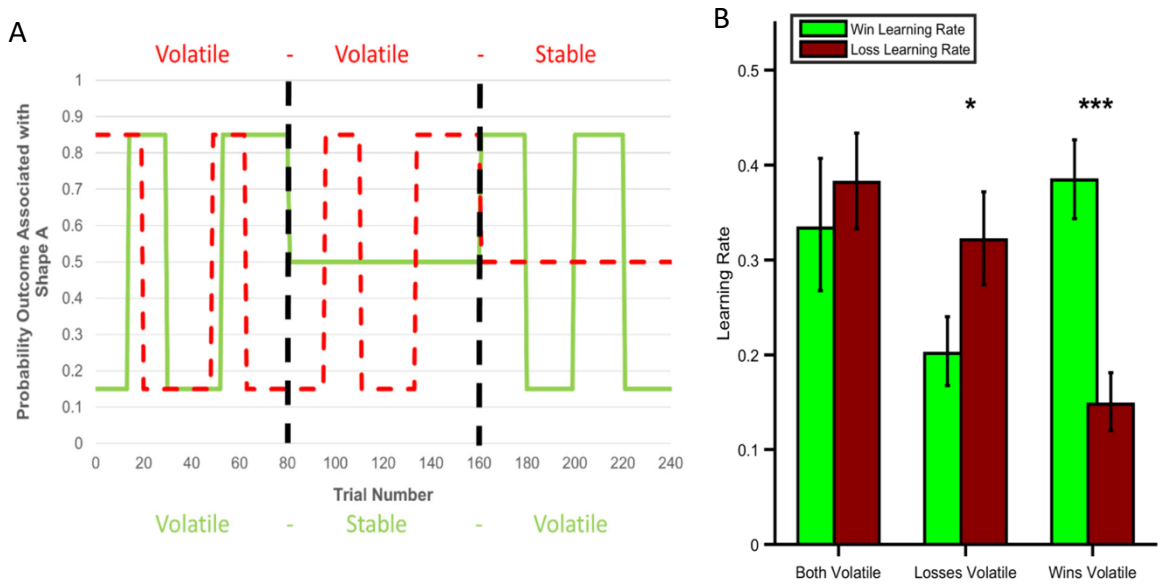

Figure 2.1. Information Bias Learning Task (IBLT). (A) Underlying reward structure for the 'both-volatile', 'loss-volatile' and 'win-volatile' conditions. In this task, the volatility of the wins and losses is manipulated independently. In 'wins-volatile' blocks, the wins are associated with one of the shapes in 75% of the trials, and with the other in 25%. This association reverses a few times within the block. Losses are randomly presented with either shape (50%) and are therefore uninformative. In 'losses-volatile' blocks, the probability pattern is reversed. In 'both-volatile' blocks, wins and losses are independently associated with one shape in 75% and with the other in 25%. (B) Healthy participants adjust their learning rates according to the volatility context, i.e. they show higher learning rates for the volatile compared to the stable outcome (losses-volatile and wins-volatile condition). Retrieved from (Pulcu & Browning, 2017).

In the first research project, two proof-of-concept studies have been conducted to test the hypothesis that tDCS applied during reinforcement learning can reduce negative biases. As discussed above, negative biases might result from increased punishment versus reward learning (Aylward et al., 2019; Eshel & Roiser, 2010; Pike & Robinson, 2022). Our main outcome of interest was therefore a relative change in punishment versus reward learning, i.e. an increase in reward learning rate, and/or decrease in punishment learning rate. The following sections will

discuss evidence supporting the hypothesis that tDCS can reduce negative biases (2.2), and explain the choice of the experimental paradigm and stimulation setup (2.3).

2.2 Can tDCS induce a positive bias in reinforcement learning?

While learning effects of tDCS are best established in the motor domain, research in the past decade has explored effects of tDCS on cognitive domains, such as memory, attention and information processing. Many studies targeted the left DLPFC, with some of them using a bifrontal montage with the anode placed over the left, and the cathode over the right DLPFC. According to one theory (Grimm et al., 2008; Hecht, 2010) depression is characterised by an imbalance between the left and right hemisphere. While the underactive left hemisphere is associated with the processing of positive, the overactive right hemisphere is supposed to primarily process negative information. By this reasoning, increasing excitability in the left DLPFC, and simultaneously decreasing excitability in the right DLPFC for bifrontal montages might therefore lead to a shift away from negative, towards positive information. The following sections will discuss whether tDCS can induce valence-specific effects (2.2.1) and modulate reinforcement learning processes (2.2.2).

2.2.1 Valence-specific effects of tDCS on cognition

Most studies on valence-specific effects of tDCS on cognition have focused on visual attention. The most basic studies analysed how long participants' gaze was fixated on negative stimuli. In one study, anodal tDCS was applied to the left DLPFC during a dual-video stressor task (Chen, Basanovic, Notebaert, MacLeod, & Clarke, 2017). During real compared to sham tDCS, participants spent less time looking at threatening stimuli in the videos. In another study measuring the speed of directing the gaze towards or away from emotional faces, anodal tDCS applied to the left DLPFC enabled participants to disengage their gaze more quickly from emotional faces (Sanchez-Lopez, Vanderhasselt, Allaert, Baeken, & De Raedt, 2018). Although the effect of tDCS

in this study was not valence-specific, this finding suggests that tDCS can increase top-down regulation of attention towards emotional information.

Two studies investigated the effect of tDCS on an attentional dot-probe task. In this task, image pairs were presented, and participants were asked to respond as quickly as possible to a probe presented behind one of the images. Stimuli were emotional pictures, i.e. happy, fearful or neutral faces. Vigilance to threat was assessed by calculating the difference in reaction time when the probe located behind the fearful face compared to when it was presented behind the neutral face. While healthy participants reacted faster towards threatening compared to neutral faces in the sham condition, bifrontal tDCS diminished this effect (Ironside et al., 2016). In another study, anodal tDCS applied to the left DLPFC had no effect on performance in a dot-probe task (Clarke et al., 2020). One possible explanation for the absence of an effect in this study might be that more complex stimuli, such as war scenes, were presented, and that the stimuli differed in arousal levels. However, tDCS reduced emotional reactivity to negative video clips in this study.

While the studies mentioned above were conducted in healthy participants, three studies investigated the effect of tDCS on affective biases in anxiety or depression. One tDCS-fMRI study examined the effect of bifrontal tDCS on attentional capture by threat in high trait anxious individuals (Ironside et al., 2019). After receiving bifrontal tDCS, participants were less distracted by threatening stimuli, which was associated with reduced response of the amygdala to threat and increased activity in cortical control areas. Another study investigated the effect of bifrontal tDCS on affective biases in participants diagnosed with depression (Brunoni et al., 2014). This study used a Stroop-task with emotional words in which individuals suffering from depression have been found to show a negative bias (Epp, Dobson, Dozois, & Frewen, 2012). Bifrontal compared to sham tDCS reduced the reaction time to negative words, i.e. participants were less distracted by negative words while receiving bifrontal tDCS. A third study investigated tDCS effects on affective biases at different processing levels in individuals diagnosed with generalised anxiety disorder (Nejati, Khalaji, Goodarzi, & Nitsche, 2021). Anodal tDCS applied to the right

ventromedial prefrontal cortex and left DLPFC reduced attentional bias towards threatening stimuli in a dot-probe task. Combined anodal DLPFC and cathodal ventromedial prefrontal cortex stimulation reduced negative interpretation bias in a task in which participants were asked to interpret images of eyes.

While the studies described above tested the effect of tDCS in paradigms *assessing* affective biases, one study investigated the effect of tDCS in a paradigm designed to *induce* affective biases (Clarke et al., 2014). This is of particular relevance since the project discussed in this chapter is motivated by the idea of combining tDCS with a learning task designed to induce a positive bias. The paradigm employed in this study (Clarke et al., 2014) was an attentional bias modification task. In the task, word pairs consisting of a neutral and a threat-related word were presented on a screen. To bias participants' attention towards or away from threat, a probe consistently replaced the threat-related or neutral word, respectively. In a sample of participants with a medium level of trait anxiety, anodal tDCS applied to the left DLPFC during task performance increased the attentional bias towards or away from threat in line with the targeted direction. In another study investigating the effects of tDCS on attentional bias modification in highly anxious individuals, no such effect was observed (Heeren, Baeken, Vanderhasselt, Philippot, & de Raedt, 2015). However, tDCS reduced the duration that participants' gaze was fixed on threatening stimuli. The divergent findings of these two studies might be caused by differences in the attentional bias modification paradigm (words vs. faces), or the stimulation procedure (cathode placed on trapezius muscle vs. supraorbital area; 1mA for 17mins vs. 2mA for 25mins).

Although the findings are somewhat heterogeneous, there is evidence suggesting that tDCS applied to the prefrontal cortex can reduce negative biases in healthy participants as well as in individuals suffering from depressive or anxiety symptoms. This has mainly been shown for attentional biases, but limited evidence is also available for higher-level cognitive biases. The study by Clarke and colleagues (2014) using attentional bias modification is of particular relevance to the current project since it combines tDCS with a training paradigm designed to induce affective biases.

Importantly, the results show that tDCS increased the training effect of the task in line with the targeted direction, i.e. it increased the bias towards or away from threat depending on the task condition. This is direct evidence supporting the central hypothesis of this project that tDCS applied during a bias learning task can increase the behavioural effects of the task.

Another central hypothesis for this project is that the cognitive state during stimulation is critical. While some of the studies mentioned above applied the stimulation before task performance (Chen et al., 2017; Ironside et al., 2019; Ironside et al., 2016; Sanchez-Lopez et al., 2018), we hypothesised that tDCS applied during task performance might be more effective in reducing negative biases than tDCS applied before task performance. The rationale for this will be discussed in more detail in section 2.2.3.

While the studies discussed in this section support the hypothesis that tDCS can reduce negative biases, the vast majority of studies have focused on attentional biases. However, the current project is based on recent findings that negative biases might be caused by deficits in reinforcement learning. The next section will therefore discuss whether tDCS can modulate reinforcement learning processes.

2.2.2 Effects of tDCS on reinforcement learning

The DLPFC is part of a network involved in reinforcement learning (Haber & Knutson, 2010). Of particular interest to this project is the question whether tDCS can modulate learning rates. DLPFC activity has been related to volatility of reward associations (Farashahi, Donahue, Hayden, Lee, & Soltani, 2019; Massi, Donahue, & Lee, 2018) which suggests that stimulating this region might induce changes in learning rates.

Two studies investigated the effect of tDCS on performance in a probabilistic learning task in which participants had to learn reward associations for three symbol pairs. Within each pair, the

probabilities of receiving a reward were reciprocal, and the pairs differed in how strongly the reward was associated with one vs. the other symbol. In one study, anodal tDCS applied to the left DLPFC increased choice randomness, but did not affect learning rates (Turi et al., 2015). In a second study, cathodal tDCS applied to the prefrontal pole decreased task performance by decreasing win-stay and lose-shift probabilities to a similar extent (Casula et al., 2017). Although no computational parameters were estimated, this pattern of findings is in line with increased choice randomness. In the same study, anodal tDCS had no effect during task performance, but decreased performance in the following test phase. In the test phase, the six symbols were presented in new pairs which had not been presented in the learning task so that participants had to transfer their estimate of reward probabilities to a new context. Anodal compared to sham tDCS reduced accuracy for trials in which symbols with a similar probability were paired.

Other studies assessed the effect of tDCS in a reversal learning task in which two stimuli were probabilistically associated with a reward, and the association reversed several times over the course of the experiment. In one study, cathodal tDCS was applied to the ventrolateral or dorsomedial prefrontal cortex (Albein-Urios et al., 2019). While tDCS applied to the dorsomedial prefrontal cortex had no effect, tDCS applied to the ventrolateral prefrontal cortex led to more reversal errors compared to sham tDCS. Since the stimulation selectively increased the number of errors right after the reversal of probabilities, but not other types of errors, the observed effect is in line with a decreased learning rate. In a second study, the reversal learning task was analysed with a computational model that separately modelled learning rates for the chosen and the unchosen option (Panitz et al., 2022). Anodal tDCS applied to the left medial prefrontal cortex increased the lose-switch probability which was related to an increased learning rate from the unchosen option.

One study investigated the effect of tDCS applied to the right frontopolar cortex on exploration vs. exploitation (Raja Beharelle, Polania, Hare, & Ruff, 2015). In the task, participants had to choose between three slot machines whose payoff values slowly drifted over time. Participants had to decide whether to keep on choosing a known option (“exploit”) or test one of the untried

alternative options which might potentially be better than the known option ("explore"). Anodal tDCS increased the number of exploratory choices, whereas cathodal tDCS decreased it. The effects of tDCS were related to increased or decreased influence of prediction errors of the highest-paying option, which is in line with an increase or decrease in learning rates from the highest-paying option.

In another study, participants had to choose between two options that differed in reward magnitude and reward probability (Hämmerer, Bonaiuto, Klein-Flügge, Bikson, & Bestmann, 2016). The reward magnitudes on each trial were explicitly presented, whereas the probabilities had to be inferred from the outcomes and drifted over the course of the experiment. The authors simulated the effect of anodal or cathodal tDCS applied to the medial prefrontal cortex in a biophysical attractor model. The simulations predicted that depolarising (or hyperpolarising) neuronal populations, which are hypothesised to represent the expected value of different options, would increase (or decrease) choice randomness, but would not affect learning rates. In line with these simulated predictions, the experimental data confirmed that anodal tDCS applied to the ventromedial PFC increased choice randomness without affecting learning rates.

Taken together, the results obtained in the few studies that have been conducted so far are heterogeneous. In some studies, anodal as well as cathodal tDCS applied to different prefrontal regions increased choice randomness but did not affect learning rates (Casula et al., 2017; Hämmerer et al., 2016; Turi et al., 2015). Since the induced current interferes with the ongoing computations performed by the stimulated brain regions, tDCS of any polarity might simply induce entropy, impairing decision-making processes which might result in increased choice randomness. Interestingly, the simulations performed in the study by Hämmerer and colleagues (Hämmerer et al., 2016) predict that cathodal tDCS should lead to a decrease in choice randomness, i.e. according to the simulations it should theoretically be possible to improve decision-making with tDCS. However, a study that applied cathodal tDCS found increased choice randomness (Casula et al., 2017).

Other studies found that anodal tDCS increased whereas cathodal tDCS decreased learning rates without affecting choice randomness (Albein-Urios et al., 2019; Panitz et al., 2022; Raja Beharelle et al., 2015). While these results might appear to contradict the studies mentioned above, an interesting pattern emerges when considering the experimental paradigms. The majority of studies that used paradigms in which the reward association changes over time (e.g. reversal learning) reported effects of tDCS on learning rates. Most studies that used paradigms with fixed reward probabilities of several choice pairs (e.g. probabilistic instrumental learning) observed tDCS effects on the inverse temperature. This suggests that tDCS affects the cognitive processes that are predominantly engaged during task performance. In tasks with stable reward associations of multiple choice-pairs, the main challenge is to track and remember the occurrence of outcomes for each choice pair; failure in this will lead to increased choice randomness. In tasks involving volatile reward associations of a single choice pair, remembering the occurrence of outcomes is easier, and the main challenge is to decide when to switch to the other option, i.e. choose the optimal learning rate. However, it should be noted that learning rates can be better estimated in tasks with volatile rather than stable reward associations so that tasks with stable reward associations might be less sensitive to changes in learning rates.

In the current project, we aimed at increasing win learning rates using tDCS. Therefore, we have chosen a paradigm which manipulates the volatility of win and loss outcomes separately, thus requiring participants to choose the optimal win and loss learning rate. In line with the theory that tDCS might specifically affect the decision-making processes most engaged in the task (see next section as well), we hypothesised that bifrontal tDCS would increase (win) learning rates. Based on the evidence on valence-specific effects of tDCS reported in the previous section, we hypothesised that the effect would be valence-specific, i.e. tDCS should selectively increase *win* learning rates. One central assumption of the project is that tDCS only causes this effect if it is applied *during* rather than *before* task performance, i.e. the cognitive state during stimulation matters. The next section will therefore discuss evidence suggesting that the effects of tDCS depend on the state of the brain.

2.2.3 Brain-state dependency of tDCS effects

This project is motivated by the idea that applying tDCS during performance of a learning task might be more effective in depression treatment than applying tDCS at rest. This rationale builds on findings that tDCS induces plasticity and increases learning effects (Fritsch et al., 2010; Reis et al., 2009). Applying tDCS during the performance of a task should strengthen the brain circuits engaged in the task, e.g. circuits relevant to reward processing. We therefore hypothesised that tDCS would increase reward learning rates or decrease punishment learning rates if applied *during* task performance ("online tDCS"), but not if applied *before* task performance ("offline tDCS").

Physiological evidence supporting this logic comes from a study investigating the effects of direct current stimulation in mouse brain slices (Fritsch et al., 2010). The application of a direct current led to long-lasting potentiation of post-synaptic potentials only if low-frequency stimulation was applied at the same time. This indicates that the effect of direct current stimulation on plasticity depends on synaptic activation, i.e. tDCS should induce plasticity specifically in the pathways that are activated during the stimulation period.

On the behavioural level, multiple studies have demonstrated that tDCS applied during task performance can increase learning effects. This is best established in the motor learning domain. For example, anodal tDCS applied to the motor cortex was found to increase offline consolidation in a motor learning task performed over five training sessions (Reis et al., 2009). Several studies showed that anodal tDCS applied to the motor cortex in a single session can increase performance in a motor task or increase the longevity of learning effects (Fritsch et al., 2010; Galea, Vazquez, Pasricha, de Xivry, & Celnik, 2011; Herzfeld et al., 2014; O'Shea et al., 2017; Reis et al., 2009).

Immediately relevant to the current project is the aforementioned study in which tDCS was applied during attentional bias modification (Clarke et al., 2014). Anodal compared to sham tDCS increased the effect of attentional bias training in the targeted direction, i.e. participants trained to

attend to threatening information did even more so when anodal tDCS was applied, and vice versa for participants who were trained to avoid threatening stimuli. This shows that tDCS did not have a general effect on task performance but specifically increased the learning effect induced by the task.

Some studies directly compared the effects of online and offline tDCS. One study on motor learning found that tDCS applied to the motor cortex during a motor-learning task had a polarity-specific effect. Anodal tDCS induced faster, and cathodal tDCS slower learning (Stagg et al., 2011). TDCS applied before task performance slowed down learning independently of the polarity. The different effects observed for stimulation during vs. before task performance confirm that tDCS effects depend on the brain state. In an imaging study, tDCS was applied to the right inferior frontal gyrus during a choice reaction time task or at rest (Li et al., 2019). While tDCS applied at rest activated the default mode network and deactivated the salience network, tDCS applied during task performance activated the salience network. Since the default mode network is active at rest, and the salience network is associated with switching from a rested to an engaged brain state (Menon, 2011), this pattern of results supports the idea that tDCS increases functional activation in the networks activated during the stimulation period.

Several studies tested whether online or offline tDCS induces stronger enhancing effects on cognition. Effects on visuo-spatial learning, cognitive training, language and memory have been investigated. Throughout all areas, results are very heterogeneous. While some studies suggest that online tDCS is superior to offline tDCS (de Aguiar, Paolazzi, & Miceli, 2015; Lee, Lee, & Kang, 2021; Martin, Liu, Alonzo, Green, & Loo, 2014; Oldrati, Colombo, & Antonietti, 2018), other studies suggest the opposite (Barbieri, Negrini, Nitsche, & Rivolta, 2016; Friehs & Frings, 2019; Grasso, Tonolli, & Miniussi, 2020; Pirulli, Fertonani, & Miniussi, 2013; Summers, Kang, & Cauraugh, 2016; Westwood & Romani, 2017).

Taken together, consensus in the field is that the effects of tDCS depend on the brain state. While evidence from the cognitive domain so far does not suggest superiority of online over offline tDCS, several indirect findings support this theory: physiologically, tDCS is hypothesised to strengthen pathways activated during the stimulation period; behaviourally, tDCS has been shown to enhance learning effects acquired during the stimulation period.

2.3 Rationale of the experimental setup

The aim of the project was to investigate whether tDCS can reduce negative biases in reinforcement learning, i.e. increase reward learning rates or decrease punishment learning rates. Learning rates are known to be influenced by the volatility of action-outcome associations (Behrens et al., 2007). We therefore chose the Information Bias Learning Task (IBLT)(Pulcu & Browning, 2017) which manipulates the volatility of positive and negative events separately. This task design makes it possible to estimate to what extent participants' choices are influenced by positive vs. negative outcomes. In different conditions of the task, one of the outcomes (either positive or negative) or both outcomes are volatile. Participants have been shown to adjust their learning rates appropriately to the volatility condition (Pulcu & Browning, 2017). This paradigm is optimal for the purpose of the project for three reasons: (1) Negative biases have been suggested to result from an overestimation of the informativeness of negative outcomes (Pulcu & Browning, 2017). The design of this paradigm allows for the assessment of potential negative biases. (2) Across the different task conditions, the paradigm manipulates which outcomes (positive or negative) should be given more weight to optimise behaviour. Therefore, it can be analysed how tDCS affects learning rates depending on the informativeness of different events. (3) Since the paradigm manipulates which type of outcome should be prioritised, for future studies, the paradigm can be transformed into a positive bias training by repeatedly presenting the condition in which positive outcomes are more informative than negative outcomes.

We chose a bifrontal tDCS montage with the anode placed over the left, and the cathode over the right DLPFC. This montage is one of the two most commonly used montages in clinical trials of depression treatment (Razza et al., 2020). A previous study has compared the effect of the two most commonly used tDCS montages, and found that bifrontal tDCS induced larger cognitive effects than a unilateral montage (with the cathode placed on the right supraorbital ridge)(Ironside et al., 2016). According to electric field modelling (Karabanov et al., 2019), the net effect of a bifrontal electrode setup appears to be an increase in excitability in the left hemisphere associated with the processing of positive information, and a decrease in excitability in the right hemisphere associated with the processing of negative information (Hecht, 2010). The DLPFC is part of a network involved in reinforcement learning and has been shown to track the volatility in the environment (Farashahi et al., 2019; Massi et al., 2018). Hence, bifrontal tDCS might counter-act negative biases in reinforcement learning which might manifest in increased learning rates from positive outcomes or decreased learning rates from negative outcomes.

This hypothesis was tested in two proof-of-concept studies in healthy individuals (chapter 4) and individuals suffering from low mood (chapter 6). The next section of the Introduction will discuss the theoretical background for the second part of the project (cognitive phenotyping).

3 Cognitive phenotyping

Depression is associated with large variability in symptomatology as well as in treatment response. It has therefore been suggested that different phenotypes of depression might exist with different underlying psychopathology. Individuals with the same phenotype of depression might be similar in psychopathology and might therefore respond to the same treatment. Past attempts at defining phenotypes of depression have focused on self-reported symptoms, and more recently on neuroimaging markers. Based on the key role that cognition plays in the development and maintenance of depressive symptoms, we propose that phenotypes of depression could be defined based on cognitive markers. With respect to NIBS treatment, identifying each individual's cognitive phenotype could help to select the best TMS target region for each individual. Different cognitive phenotypes might best respond to TMS applied to the brain regions the underlying cognitive dysfunctions are associated with. Previous research indicates that the relative severity of anxiety vs. anhedonia might be relevant for distinguishing between different phenotypes. The aim of this project was therefore to investigate whether and which parameters extracted from five different cognitive tasks could capture distinct symptom dimensions related to anxiety and anhedonia. Relevant parameters could be included in future clinical trials to test if they have the potential to guide treatment selection, including (but not limited to) the choice of target region for TMS treatment.

3.1 Personalising NIBS treatment

One of the biggest challenges in depression treatment is finding the right treatment for the right person. There is large heterogeneity in treatment response, and there is currently no procedure to predict who will respond to which treatment. In addition, depression is associated with a large variety of different symptoms. Patients with the same diagnosis ("depression") might have very little overlap in their symptoms. The variability in symptomatology and treatment response suggests that multiple different phenotypes of depression with different underlying

psychopathology might exist (Beijers, Wardenaar, van Loo, & Schoevers, 2019). A potential solution could be to tailor treatment to individuals, informed by phenotyping. Treatment selection or stratification could be guided by identifying each patient's depression phenotype since patients with the same phenotype might respond to the same treatment.

While in most clinical trials the DLPFC has been stimulated, research on the mechanisms of NIBS treatment highlights that multiple alternative candidate target regions exist (see 1.3.2). Attempts at personalising TMS treatment have therefore focused on personalising the target region. Increasing evidence suggests that patients differ in the target region they best respond to (Drysdale et al., 2017; Fox et al., 2012; Siddiqi et al., 2020b). For example, there is evidence that patients who predominantly report symptoms of anxiety might respond better to TMS applied to the dorsomedial prefrontal cortex, whereas patients suffering from higher anhedonia might respond better to TMS applied to the ventromedial prefrontal cortex (Downar et al., 2014). In recent literature, two general approaches to personalising the target region have emerged. One approach is to optimise the stimulation target locally within DLPFC to take account of individual differences in structural and functional anatomy. Another is to categorise patients into different subgroups and tailor stimulation differently for each group by targeting distinct cortical regions.

3.1.1 Adjusting the target region based on individual anatomy

The influential study by Fox and colleagues (Fox et al., 2012) found that clinical improvement in response to TMS treatment was related to the strength of resting state functional connectivity between the DLPFC stimulation site and the sgACC. This work suggests that TMS treatment could be personalised by stimulating in each individual that subregion within the DLPFC with the strongest anticorrelation to the sgACC. However, the study used group-averaged connectivity, precluding direct conclusions about the relationship between individual connectivity and treatment response. Two studies have tried to replicate the findings on an individual-patient level. One study (Cash et al., 2019) found that treatment outcome could be predicted based on individual

connectivity between the DLPFC and sgACC. In another study (Siddiqi, Weigand, Cooke, Pascual-Leone, & Fox, 2019), no such relationship was found. Therefore, it is currently unclear whether this approach to personalisation is likely to enhance treatment effectiveness. Translating neuroimaging analyses based on group averages to individuals is often difficult due to methodological problems such as low signal-to-noise ratio in the areas of interest or low test-retest reliability of functional connectivity measures.

A point that highlights that treatment individualisation is still in early stages is the fact that all of the above-mentioned studies used retrospective analyses. Only one clinical trial has attempted to personalise the stimulation site in advance of treatment. In the Stanford Neuromodulation Therapy trial (Cole et al., 2022) each patient's subregion within the DLPFC most anticorrelated with the sgACC was targeted. Personalisation of the stimulation site was combined with a new accelerated high-dose intermittent theta-burst protocol. In a sham-controlled trial, a response rate of 85% and remission rate of 78% were observed. However, the trial did not include any control groups that could provide insights into the extent to which the personalised stimulation site contributed to the increased response rate.

3.1.2 Categorisation into different subtypes of depression

An alternative approach is to categorise patients into different depression subtypes and target the brain regions thought to best engage the relevant dysfunctional brain circuitry. One study used this approach to retrospectively investigate potential predictors of treatment response to dorsomedial prefrontal cortex TMS (Downar et al., 2014). Non-responders showed significantly stronger anhedonia at baseline, and lower connectivity in a reward pathway comprising the ventral tegmental area, the striatum and lateral orbitofrontal cortex. The authors suggested the existence of two different depression subtypes, one with preserved hedonic function that responds to dorsomedial prefrontal cortex TMS, and one with pronounced anhedonia that is unresponsive to

dorsomedial prefrontal TMS. This latter group had abnormal connectivity in lateral orbitofrontal cortex, suggesting that region as a potential alternative personalised TMS treatment target.

In that study, the categorisation into subtypes was based on the clinical response to TMS treatment. In contrast, a recent study hypothesised that depressed patients could be grouped into subtypes based on correlated clinical and imaging markers (Drysdale et al., 2017). The authors used canonical correlation analysis to extract two dimensions based on correlated symptoms and resting-state connectivity. One dimension was related to fronto-striatal and orbitofrontal connectivity markers which correlated with anhedonia and psychomotor retardation. The second dimension was related to connectivity in limbic and lateral prefrontal areas which correlated with anxiety and insomnia. A clustering algorithm applied to these two dimensions suggested the existence of four different "biotypes" of depression characterised by different patterns of functional connectivity that correlated with differing levels of anhedonia and anxiety symptoms (Figure 3.1A). These four biotypes also differed in their responsiveness to TMS treatment applied to the dorsomedial prefrontal cortex (Figure 3.1B), which suggested that the different biotypes may have different optimal TMS treatment targets. Because of the potential clinical implications, this study has attracted a lot of attention, but needs to be interpreted with caution since it has also earned criticism regarding methodological shortcomings (Dinga, Schmaal, & Marquand, 2020; Dinga et al., 2019).

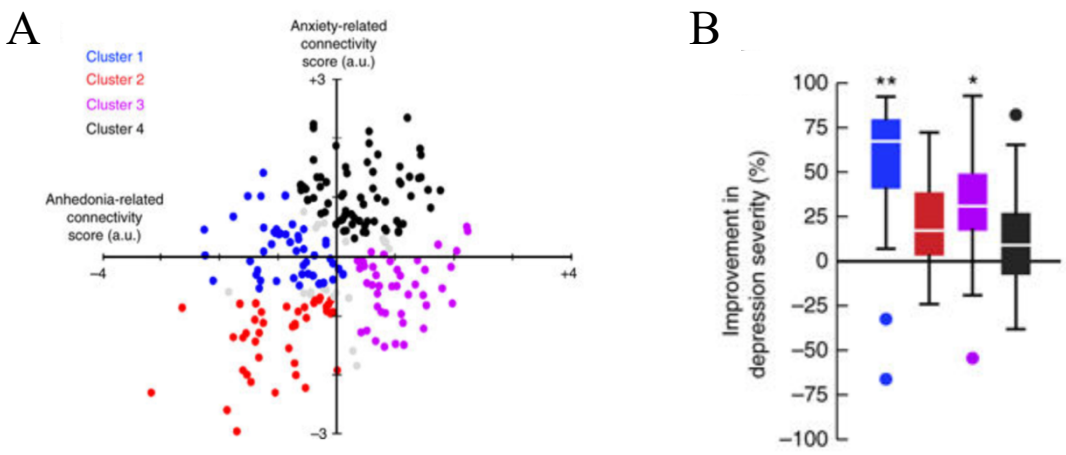

Figure 3.1. The four different biotypes of depression identified by Drysdale and colleagues based on connectivity profiles (Drysdale et al., 2017). (A) The clusters of the four biotypes plotted along the two connectivity scores. (B) Percent of improvement in depression score in response to TMS. Reproduced from (Drysdale et al., 2017) with permission.

While the previous studies aimed at identifying subcategories of depression, another study investigated how target regions are associated with improvement in specific symptom clusters (Siddiqi et al., 2020a). The authors created connectivity maps for the TMS targets of individual patients based on the resting-state data of a large connectome database. Across patients, each voxel's connectivity to the target region was correlated to improvement in each depressive symptom, so that each of the resulting maps corresponded to the degree to which each voxel's connectivity with the target region predicted improvement in a specific symptom. Since there were similarities between these symptom-response maps they were categorised into two clusters. The "dysphoric" cluster included symptoms such as sadness, decreased interest and suicidality, whereas the "anxiosomatic" cluster was associated with irritability, sexual disinterest and insomnia (Figure 3.2). In further analysis steps, symptom-response maps for the two clusters were combined, which resulted in a map indicating to what extent stimulating a region would result in the reduction of "dysphoric" or "anxiosomatic" symptoms. According to this map, stimulating the DLPFC would primarily reduce dysphoric symptoms, whereas stimulating the dorsomedial prefrontal cortex would reduce anxiosomatic symptoms. This map retrospectively explained improvements in different symptoms across fourteen different TMS studies. The results of this study have potential clinical value since the map could be used to personalise the target region based on the symptom

profile of a given patient. This needs to be validated in a prospective clinical trial to test whether such an individually tailored treatment approach does improve outcomes.

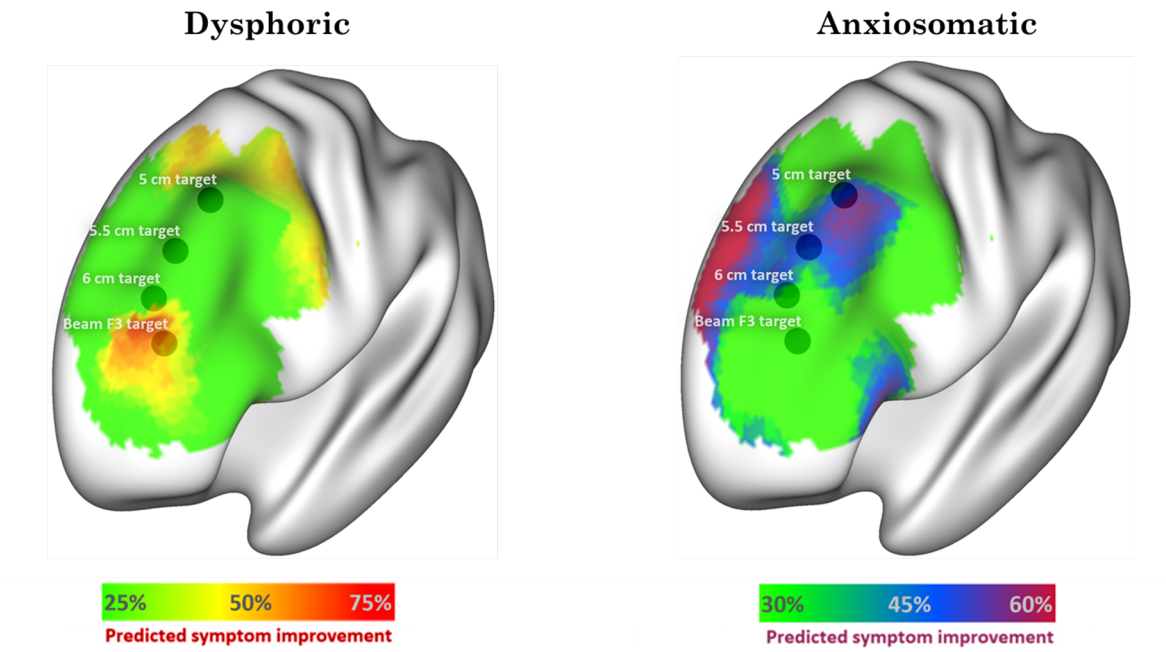

Figure 3.2. Association between potential target regions and improvement in two different symptom clusters as proposed by Siddiqi and colleagues (Siddiqi et al., 2020a). The two maps the predicted improvement in dysphoric or anxiosomatic symptoms for potential TMS target regions. The overlaid dots represent common target areas. Figure provided by Shan Siddiqi, adapted from (Siddiqi et al., 2020a).

Taken together, these three studies suggest that different patients characterised by different patterns of symptoms and brain network dysfunction might require different TMS target regions. Although the methodological approach differed between the studies, they all suggest that the relative severity of anhedonia vs. anxiety symptoms might be relevant in determining the optimal target region. The study by Downar and colleagues (Downar et al., 2014) suggests that patients with preserved hedonic functions respond better to TMS applied to the dorsomedial prefrontal cortex than patients with severe anhedonia. This is compatible with the finding by Siddiqi and colleagues (Siddiqi et al., 2020b) suggesting that stimulating the dorsomedial prefrontal cortex reduces anxiety rather than anhedonia.

3.2 Cognitive phenotyping

Previous attempts of personalising TMS treatment were based on symptom profiles or neuroimaging markers. The identification of phenotypes based on clinical symptoms has not been found to be predictive of treatment response to different pharmacological treatments (Arnow et al., 2015; Uher et al., 2011). As discussed above, recent studies have tried to identify biological subtypes based on imaging markers with the hypothesis that brain characteristics might be better in guiding treatment selection than subjective reports of symptoms (Downar et al., 2014; Drysdale et al., 2017). However, neuroimaging is not part of standard psychiatric diagnostics, limiting the practical clinical potential of this approach.

Recent research indicates that phenotypes with similar psychopathology might be identified based on cognitive functions (Gillan, Kosinski, Whelan, Phelps, & Daw, 2016; Rouault, Seow, Gillan, & Fleming, 2018). Depression and anxiety have been associated with multiple cognitive dysfunctions, especially related to reward learning and optimising behaviour in response to uncertainty (Browning et al., 2015; Gagne et al., 2020; Kube et al., 2020). The advantage of using cognitive performance rather than only clinical symptoms is that many cognitive functions have a relatively clear biological basis (e.g. goal-directed behaviour (Dolan & Dayan, 2013)). Therefore, cognitive phenotypes might be more closely associated with the underlying pathophysiology, which might improve treatment selection (Gillan et al., 2016). The use of cognitive tasks also comes with practical advantages: costs for cognitive task performance are minimal, and tasks could be performed online which makes this approach accessible to most patients. However, the practical clinical value of cognitive measures has not been investigated yet.

A variety of different cognitive task parameters have been found to be correlated with psychiatric symptoms. However, it is unclear how these different task parameters relate to one another and if they capture distinct cognitive processes affected in depression. As discussed above, previous research suggests that the relative severity of anxiety vs. anhedonia might distinguish different

phenotypes of depression. The aim of this project was therefore to test whether parameters derived from a selection of five cognitive tasks differentially relate to symptoms of anxiety or anhedonia. The results of this project will shed light on which tasks and parameters are best able to explain variance in depressive symptomatology. These tasks could then be included in future clinical trials to test whether they have potential to predict treatment response. Similar to (Drysdale et al., 2017), symptom dimensions correlated with cognitive task parameters could be used to investigate potential cognitive phenotypes of depression.

3.3 Research based on online samples with varying symptom levels

Our approach of relating task parameters to symptoms of anxiety and anhedonia builds upon previous work investigating the relationship between symptoms and task parameters in large-scale online samples with naturally varying degrees of symptoms. The current diagnostic system for psychiatry, the DSM-5, groups mental disorders into separate categories based on symptoms. However, individuals with the same diagnosis can have very different, sometimes even opposite symptoms, e.g. depression is associated with both weight gain or loss. At the same time, different diagnoses overlap in their symptoms, e.g. apathy and delusions can be associated with both depression and schizophrenia. Most importantly, diagnoses often fail to predict treatment response, likely because they do not capture the fundamental underlying mechanism causing mental illness (Insel et al., 2010).

It has therefore been argued that mental illness would be better described by multiple continuous dimensions that capture the underlying cognitive dysfunctions and cut across conventional diagnostic boundaries (Heinz, Schlagenhauf, Beck, & Wackerhagen, 2016). Recent studies have recruited large online samples with natural variations in mental health with the aim to extract latent cognitive dimensions correlated with psychiatric symptoms. One symptom dimension that has gained increasing interest is related to symptoms of compulsivity and intrusive thoughts and has been linked to deficits in goal-directed control (Gillan et al., 2020; Gillan et al., 2016; Voon,

Reiter, Sebold, & Groman, 2017), as well as increased confidence and impaired metacognition (Hoven, Denys, Rouault, Luigjes, & van Holst, 2022; Rouault et al., 2018; Seow & Gillan, 2020). A symptom dimension related to depression and anxiety has been associated with altered reward anticipation (Hägele et al., 2015), negative biases in a perceptual decision-making task (Daniel-Watanabe, McLaughlin, Gormley, & Robinson, 2020), lower confidence and higher meta-cognitive efficiency (Hoven et al., 2022; Rouault et al., 2018). Although an increasing number of studies is investigating relationships between psychiatric symptoms and cognitive task parameters, to date, there is no evidence for predictive value from clinical trials.

In this project, we are investigating the relationship of task parameters and psychiatric symptoms in a dataset collected online. Online recruitment allows for testing a much larger number of participants than could be tested in a laboratory in a similar timeframe, thus maximising statistical power. A disadvantage of online studies is that there is very little control over participants' behaviour during the online testing sessions. However, typical psychological effects have been replicated in online samples suggesting that data collected in online studies seem to be generally valid (Crump, McDonnell, & Gureckis, 2013).

Since our main goal was to identify task parameters distinguishing between symptoms of anxiety and anhedonia, we decided to recruit participants in such a way as to maximise variance in the relative severity of anxiety and anhedonia. This was achieved by screening participants using anxiety and anhedonia questionnaires, based on which participants were invited to take part in the actual study procedure. The goal of this approach was to maximise power for detecting relationships between the task parameters and symptoms of anxiety or anhedonia, in order to test which task parameters are most useful in explaining variance in symptoms. The disadvantage of our online study design is that it remains undetermined whether any findings would generalise to actual patient samples and in-person laboratory testing. Once the most informative task parameters have been identified in the online dataset, these could later be included in future clinical trials to test for predictive and explanatory potential in patient trials.

3.4 Task selection

Our project was motivated by previous work investigating approaches to personalising the TMS target region. One potential approach for task selection would therefore be to select tasks that have been shown to engage different potential TMS target regions. Individuals with deficits in a particular task might best respond to TMS applied to the key brain region engaged during that task. However, most TMS trials have focused on the DLPFC, and only very little research has explored other target regions. Therefore, it is unclear which alternative target regions would be most suitable for such an approach.

Instead, we have selected tasks that have been shown to correlate with depressive symptoms. We have chosen two paradigms that are very well established in psychiatric research (*Facial Expression Recognition Task* and *Probabilistic Instrumental Learning Task*). Although these tasks have been applied in several clinical trials, it is unclear which symptom dimension they might be most sensitive to. In addition, we have selected three more recently developed paradigms that are hypothesised to capture cognitive processes more specifically related to anxiety or anhedonia.

Facial Expression Recognition Task (FERT)

The FERT was chosen as a well-established paradigm that assesses the processing of emotional faces (Young et al., 1997). Stimuli are emotional faces that show one out of six emotions at different intensities. Participants need to decide as quickly as possible which emotion has been shown. Depression has been associated with a response bias towards negative vs. positive emotions (Gur et al., 1992; Walsh, Browning, Drevets, Furey, & Harmer, 2018; Walsh, Huneke, et al., 2018). Antidepressant drugs have been shown modulate emotional face processing (Bhagwagar, Cowen, Goodwin, & Harmer, 2004; Harmer et al., 2003; Walsh, Browning, et al., 2018), and a reduction in bias towards negative faces has been found to predict drug treatment response (Browning et al., 2019).

Probabilistic Instrumental Learning Task (PILT)

The PILT has been chosen as another well-established paradigm which assesses reward and punishment learning (Pessiglione, Seymour, Flandin, Dolan, & Frith, 2006). On each trial, one out of two symbol pairs is presented, with one pair associated with a win (win 20 points or no change), and the other associated with a loss (lose 20 points or no change). Participants are asked to maximise their winnings and are supposed to learn over time which of the symbols in a pair is more strongly associated with the outcome (win or loss). Depression has been associated with difficulties in reward learning (Kumar et al., 2018; Walsh, Browning, et al., 2018) and antidepressants have been found to normalise this deficit (Walsh, Browning, et al., 2018).

Volatility Task

The Volatility task (Browning et al., 2015) is a relatively new paradigm which assesses how learning behaviour is adjusted to the volatility of outcome associations. In the punishment learning version of this paradigm, participants are asked to choose between two different stimuli containing a number indicating the number of points that might be lost if the stimulus is chosen. The stimuli differ in the probability of actually leading to a loss. The loss probabilities are kept constant in the stable condition, and change over time in the volatile condition. Internalising psychopathology has been associated with deficits in adjusting learning rates to volatility (Browning et al., 2015; Gagne et al., 2020), with some evidence indicating that this deficit might be specific to anxiety (Browning et al., 2015).

Wheel of Fortune Task

The Wheel of Fortune Task is a new paradigm which has been developed as an ecologically valid paradigm to investigate sequential decision-making (Kolling, Scholl, Chekroud, Trier, & Rushworth, 2018). Participants are presented with a wheel of fortune which is divided into different segments. A number in each segment indicates their value and the size of the segment indicates their probability of being drawn. On each trial, participants need to decide if they would like to spin the wheel of fortune or keep the current offer and move on to the next trial. The task

requires the participant to take into account the current offer, the value and probability of the alternatives, and the cost of spinning the wheel. A recent study shows that apathy is related to a parameter capturing decision inertia, i.e. continue to search again for longer than appropriate given the value of the available options (Scholl, Trier, Rushworth, & Kolling, 2022).

Foraging Fish Game

The Foraging Fish Game is a novel task that has been developed as an ecologically valid paradigm to investigate the trade-off between foraging for rewards and vigilance to threat. Participants play a fish in the ocean and need to decide whether to forage for food, act vigilantly to watch out for predators or hide from predators. This paradigm is currently being validated. Preliminary unpublished data suggest that some parameters might specifically be related to anxiety, whereas other parameters might be related to anhedonia and apathy.

For this thesis, an exploratory dataset has been collected to examine relationships between psychiatric symptoms and the five selected cognitive tasks. Building on previous evidence suggesting that the relative severity of anxiety and anhedonia might have value for predicting treatment response, chapter 7 presents preliminary analyses testing which task parameters differentially relate to symptoms of anxiety and anhedonia.

3.5 Thesis outline

The aim of the first project (chapters 4-6) was to test whether tDCS applied during reinforcement learning can reduce negative biases. Chapter 4 examines the effect of tDCS on reinforcement learning in healthy individuals. Chapter 5 compares reinforcement learning between individuals with low mood and healthy individuals to assess how reinforcement learning is altered in low mood. Chapter 6 investigates the effect of tDCS on reinforcement learning in individuals with low mood.

The second project set out to test which cognitive task parameters might be related to different dimensions of depressive symptoms and might therefore have value for predicting treatment response. Chapter 7 investigates which task parameters derived from five selected cognitive tasks differentially relate to symptoms of anxiety and anhedonia.

4 Modulating reward learning with bifrontal tDCS in healthy adults

Transcranial direct current stimulation (tDCS) applied to the dorsolateral prefrontal cortex (DLPFC) is under investigation as an antidepressant treatment and has been found to have moderate antidepressant effects. Little is known about the cognitive mechanisms underlying the antidepressant effect of tDCS. Previous research suggests the hypothesis that, similar to other antidepressant treatments, tDCS might act by reducing negative biases. Individuals suffering from depression tend to prioritise the processing of negative over positive outcomes which is hypothesised to cause depressive symptoms. In this proof-of-concept study in healthy individuals, we investigated whether tDCS applied during a reinforcement learning task increases processing of positive information. Participants performed a reinforcement learning paradigm in which the relative informativeness of positive and negative outcomes was manipulated. Choice behaviour was analysed using a computational model with separate learning rates for win and loss outcomes. We hypothesised that bifrontal tDCS would increase processing of positive information which should lead to increased win learning rates. As hypothesised, bifrontal tDCS applied during task performance selectively increased win learning rates. This effect was replicated in an independent sample. The effect of tDCS was brain-state dependent and anatomically specific: bifrontal tDCS applied before task performance or tDCS applied to the motor cortex had no valence-specific effects. Our findings support the hypothesis that bifrontal tDCS increases processing of positive vs. negative information which might be a potential mechanism underlying the antidepressant effects of tDCS. Moreover, our results suggest that applying tDCS during the performance of a reinforcement learning task might be more effective than applying tDCS at rest.

4.1 Introduction

Transcranial direct current stimulation (tDCS) is under investigation as a novel antidepressant treatment. TDCS applied to the dorsolateral prefrontal cortex (DLPFC) has been proposed as a safer, cheaper and more accessible alternative to transcranial magnetic stimulation, which is an FDA-approved treatment. However, a recent meta-analysis showed that the antidepressant effects of tDCS are only mild to moderate (Razza et al., 2020). Therefore, further research into the mechanisms of the antidepressant effects of tDCS is needed to improve its application.

Little is known about the cognitive mechanisms of tDCS treatment. Bifrontal tDCS is hypothesised to normalise an asymmetry between the left and right hemisphere by exciting the left hemisphere associated with processing of positive information, and inhibiting the right hemisphere associated with processing of negative information (Grimm et al., 2008; Hecht, 2010). The reversal of negative biases is thought to be one mechanism underlying effective antidepressant treatments (Godlewska et al., 2016; Godlewska, Norbury, Selvaraj, Cowen, & Harmer, 2012; Harmer, Goodwin, & Cowen, 2009). There is evidence that bifrontal tDCS might also reduce negative biases. For instance, bifrontal tDCS has been shown to reduce attention to threatening faces (Ironside et al., 2019; Ironside et al., 2016) and negative words (Brunoni et al., 2014). This suggests the hypothesis that the antidepressant effect of bifrontal tDCS might reduce depressive symptoms by reversing negative biases.

Recent research in the field of computational psychiatry suggests that negative biases in depression might result from altered decision-making processes. More precisely, it has been proposed that negative biases might develop if an individual considers negative events to be more informative than positive events (Pulcu & Browning, 2017) which might lead to increased learning from negative outcomes (Aylward et al., 2019; Eshel & Roiser, 2010; Pike & Robinson, 2022).

The aim of this proof-of-concept study was to test whether bifrontal tDCS induces a positive bias in reinforcement learning in healthy participants. Bifrontal tDCS was applied during the performance of the Information Bias Learning Task (IBLT)(Pulcu & Browning, 2017). By varying the informativeness of positive and negative outcomes separately, the task manipulates whether more weight should be given to positive or negative outcomes. Thus, this task allows to test whether bifrontal tDCS modulates learning rates in a valence-specific way depending on the relative informativeness of positive and negative outcomes.

The DLPFC is part of a network involved in reinforcement learning (Haber & Knutson, 2010). DLPFC activity has been related to volatility of reward associations (Farashahi et al., 2019; Massi et al., 2018) suggesting that tDCS applied to the DLPFC might modulate learning rates. TDCS applied to the prefrontal cortex has indeed been found to alter learning rates: anodal tDCS increased (Panitz et al., 2022; Raja Beharelle et al., 2015) and cathodal tDCS decreased learning rates (Albein-Urios et al., 2019; Raja Beharelle et al., 2015) in reinforcement learning paradigms with volatile reward associations.

Based on the evidence that bifrontal tDCS might reduce attention to negative stimuli (Brunoni et al., 2014; Ironside et al., 2019; Ironside et al., 2016) and modulate learning rates in reinforcement learning tasks (Albein-Urios et al., 2019; Panitz et al., 2022; Raja Beharelle et al., 2015), we hypothesised that (1) bifrontal tDCS applied during the IBLT would increase the learning rate from positive outcomes and/or decrease the learning rate from negative outcomes. TDCS has been associated with neuroplasticity and learning effects which depend on the cognitive state during the stimulation period (Fritsch et al., 2010; Li et al., 2019; Stagg et al., 2011). We therefore hypothesised that (2) the cognitive state during stimulation is critical, i.e. tDCS applied before task performance should not increase learning from positive outcomes. Moreover, (3) the effect of tDCS should be anatomically specific, i.e. tDCS applied to the primary motor cortex as anatomical control region should not increase learning from positive outcomes or decrease learning from negative outcomes.

4.2 Methods

4.2.1 Sample

80 healthy young participants[2] (45 women, mean age = 24.71, SD ± 5.08, see Table 4.1) took part in the study. Participants were excluded from the study if they had a diagnosis of a psychiatric or neurological disease or met any contraindications for tDCS, such as a family history of epilepsy or current medication. The study has been approved by the University of Oxford Central University Ethics Committee (RE48995/RE002) and all participants gave informed written consent before taking part in the study.

Table 4.1. Mean (SD) baseline characteristics across studies.

	Online DLPFC (n = 20)	*Offline DLPFC (n = 20)*	*Online M1 (n = 20)*	*Online DLPFC (replication) (n = 20)*
Sociodemographic data				
Female (%)	9 (45.0)	15 (75.0)	14 (70.0)	7 (35.0)
Age, years	25.0 (4.3)	25.2 (5.8)	24.6 (6.4)	24.2 (3.6)
Clinical measures				
STAI-Trait	37.5 (8.9)	38.4 (7.9)	32.6 (8.7)	35.6 (5.9)
BDI	5.1 (5.2)	5.6 (8.0)	3.1 (3.4)	4.0 (3.8)

BDI: Beck Depression Inventory-II, score range = 0-63; STAI-Trait: State-Trait Anxiety Inventory (trait form), score range = 20-80.

4.2.2 Experimental design

Participants were assigned to one out of four tDCS conditions. 20 participants received bifrontal tDCS during task performance ("online tDCS" group). To test whether the observed effects were reproducible in an independent sample, another 20 participants were recruited and underwent the same tDCS protocol ("replication" group). To investigate our hypothesis that the effect of tDCS is

[2] The author of this thesis collected data for 20 participants ("M1" group). This study has been published as a journal article on BioRxiv (Overman, Sarrazin, Browning, & O'Shea, 2021). All analyses and figures reported in this chapter were re-created independently by the author, with permission from all co-authors.

cognitive-state dependent, 20 participants received tDCS at rest ("offline tDCS" group) and performed the task after the stimulation has ended. To demonstrate that the effect of tDCS is anatomically specific, another 20 participants received tDCS applied to the primary motor cortex while performing the task ("M1" group). All participants took part in two testing sessions in which they received either real or sham tDCS, the order of which was counter-balanced across participants.

4.2.3 Questionnaires

To assess the level of symptoms of depression and anxiety, the Beck Depression Inventory (BDI-II)(Beck et al., 1996) and the trait-anxiety scale from the State-Trait Anxiety Inventory (STAI)(Spielberger, 1983) were completed at the beginning of the first session. As expected, participants had low scores in both questionnaires (Table 4.1), i.e. low levels of depressive symptoms and anxiety. To test for potential changes in affect over the course of the testing sessions, the Positive and Negative Affect Schedule (PANAS)(Watson, Clark, & Tellegen, 1988) and the state anxiety scale from the STAI were completed before and after task performance. Analyses of the questionnaire data are included in the supplementary material.

4.2.4 Experimental Paradigm

Participants performed the Information Bias Learning Task (IBLT) developed by Pulcu and Browning (2017)(Figure 4.1). On each trial, a fixation cross was presented in the middle of the screen and two abstract shapes (letters from the Agathodaimon font) on the left and right side. Participants were asked to choose one of the two shapes by pressing one of two keys on a keyboard. After participants indicated their choice, a win and a loss outcome appeared on the screen. The win and a loss outcomes were independently associated with one of the two shapes, which resulted in one of four different scenarios: Participants received either the win, the loss, both outcomes or neither outcome. At the start of the task, the total amount was set to £1.50. The win

and loss outcomes were associated with an increment or decrement of 10p. If participants received both outcomes or neither outcome, the total was unchanged on that trial.

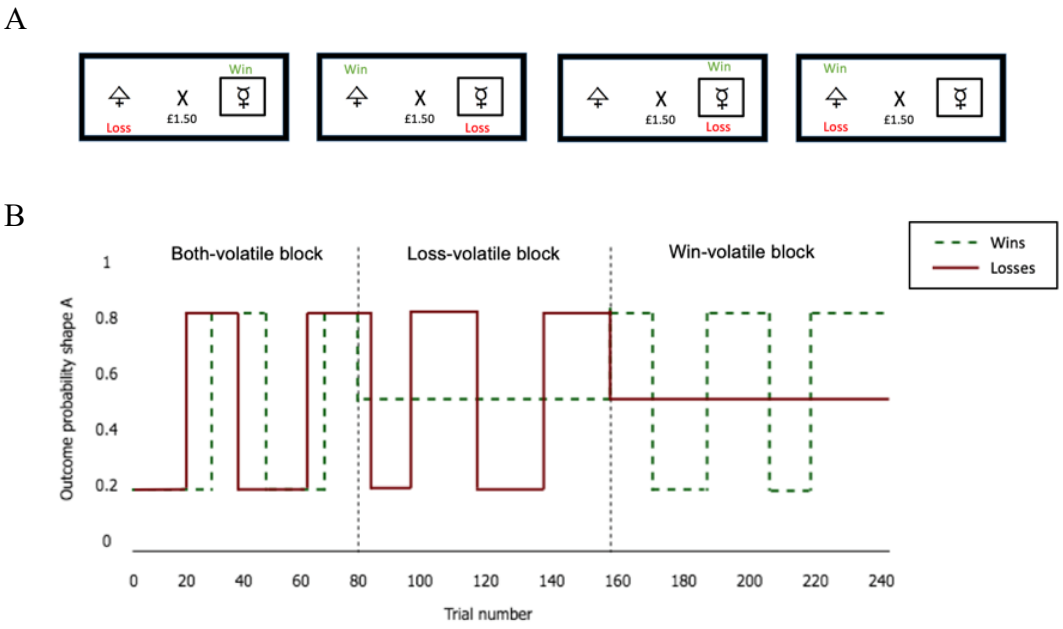

Figure 4.1. Task design of the Information Bias Learning Task. (A) On each trial, participants are asked to choose one of two shapes, by pressing the keys 'A' or 'L' for the left or right shape, respectively. Subsequently, a win and a loss outcome appeared on the screen. The win and loss outcomes are independent of each other, resulting in four possible scenarios: The chosen shape might be associated with the win, the loss, both outcomes or neither. Wins and losses were associated with an actual win or loss of 10p on each trial, respectively. (B) Underlying reward structure for a 'both-volatile', 'loss-volatile' and 'win-volatile' block. In this task, the volatility of the wins and losses was manipulated independently. In 'wins-volatile' blocks, the wins were associated with one of the shapes in 80% of the trials, and with the other in 20%. This association reversed a few times within the block. Losses were randomly presented with either shape (50%) and are therefore uninformative. In 'losses-volatile' blocks, the probability pattern was reversed. In 'both-volatile' blocks, both wins and losses are associated with one shape in 80% and with the other in 20% of the trials but are still independent of each other. Adapted from (Overman et al., 2021).

In this task, the volatility of win and loss outcomes was manipulated separately. A volatile outcome was associated with one of the shapes in 80% of the trials, and with the other in 20%. This association was volatile and reversed a few times within a task block. A stable outcome appeared randomly with each shape in 50% of the trials. In "both-volatile" blocks, both win and loss outcomes were volatile. In "wins-volatile" blocks, wins were volatile while losses were stable, which was reversed in "losses-volatile" blocks. The task consisted of six blocks of 80 trials (Figure 4.2). The task started with a both-volatile block as a measure of baseline performance and potential

bias towards learning from wins or losses. Participants then performed two wins-volatile and two losses-volatile blocks in alternating order. Whether the wins-volatile or losses-volatile block was presented first was counter-balanced across participants. The task ended with another both-volatile block.

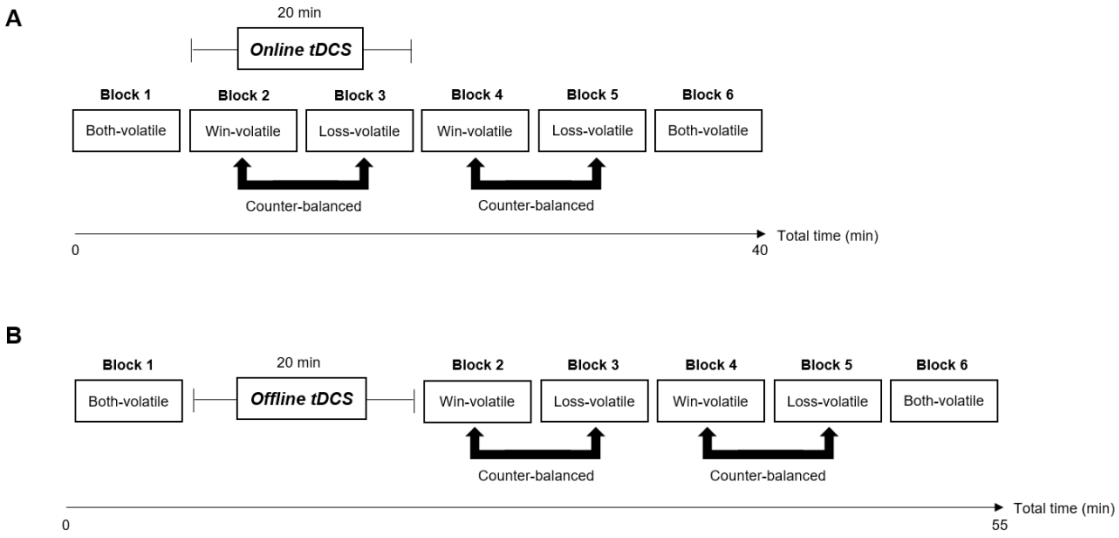

Figure 4.2. Illustration of the study design. Participants started with a "both-volatile" block, and then underwent two "wins-volatile" and two "losses-volatile" block in alternating order. Half of the participants performed the "wins-volatile" block first, while the other half performed the "losses-volatile" block first. The experiment ended with another "both-volatile" block. "Online" tDCS (participants in the "online", "replication" and "M1" groups) was applied during the second and third block, whereas "offline" tDCS was applied between first and second block at rest. Retrieved from (Overman et al., 2021).

4.2.5 tDCS protocol

The stimulation was delivered through a battery-powered DC-Stimulator Plus (neuroConn). All participants took part in two testing sessions in which they received real or sham tDCS in counter-balanced order. For both types of tDCS, the protocol included a ramp-up and ramp-down period of 10s. Real tDCS was applied for 20 minutes at an intensity of 2mA. For sham stimulation, tDCS was applied at 2mA for 30 seconds, after which the current was ramped down. During the remaining stimulation period, short pulses (110 μA over 15ms) were delivered every 550ms which does not have any physiological effects but enables impedance control. The type of stimulation delivered was determined by a code that was typed into the device to enable double-blinding.

At the start of the tDCS setup, head measurements were taken to locate the vertex. For the "online", "offline" and "replication" group, an EEG cap was used to mark the F3 and F4 electrode position (10-20 system) which are an approximation for the left and right DPLFC. The anode was placed on the left, and the cathode on the right DLPFC. In the "M1" group, the motor cortex was localised by measuring 5cm lateral from the vertex. The anode was placed on the left, and the cathode on the right motor cortex. The marked areas of the scalp were cleaned using saline solution. tDCS was delivered via 5x5cm^2 rubber electrodes placed into saline-soaked sponge cases which were fixed on the scalp using two elastic bands. The impedance was kept below 10μΩ.

4.2.6 Experimental procedure

After providing informed written consent and completing the questionnaires, the tDCS equipment was set up. The stimulation was turned on for a few seconds to make sure the participant was comfortable and the impedance was below 10μΩ. Afterwards, the task was explained to the participant. The participant was instructed that the goal in the task was to make as much money as possible. The participant was told that the wins or losses were associated with one of the two shapes but that the association was not 100% so that they would sometimes appear with the other shape. They were told that they could learn over time which shape the wins or losses were associated with, and that the association could change over time. They were also told that the side the shapes were presented on was irrelevant to the task. Participants then completed 10 practice trials to make sure they understood the instructions. The instructions and practice trials were repeated if necessary.

After performing the first task block (both-volatile), the stimulation was turned on for 20 minutes. Participants in the "online", "replication" and "M1" groups performed the second and third task block during the stimulation period. Participants were asked to pause after the third task block and wait for the stimulation to finish which took around five minutes. After the stimulation had

finished, participants performed the remaining three task blocks. Participants in the "offline" group were stimulated at rest, and completed the remaining five task blocks afterwards.
After task completion, the tDCS setup was removed, and participants were asked to complete another set of questionnaires. This included the state anxiety scale of the STAI and the PANAS to measure potential changes in mood. A questionnaire about adverse effects was used to assess the sensations participants were experiencing during the stimulation period, and whether they believed that they had experienced real or sham tDCS. A final questionnaire asked participants about strategies they used during the task.

4.2.7 Computational modelling

Performance in the IBLT was analysed using a computational model that was fit to the participants' trial-by-trial choices. Six different models were fit separately to each participant's choices in each block. The fit of these six models was compared using the Bayesian Information Criterion (BIC) averaged across all participants, blocks and condition.

Model 1. The winning model used a modified version of a Rescorla-Wagner updating rule in which the probability of an outcome being associated with shape A was modelled separately for win and loss outcomes. The separate probabilities of a win or loss being associated with shape A were modelled with separate learning rates for wins and losses:

$$rwin_{(i+1)} = rwin_{(i)} + \alpha win * (winout_{(i)} - rwin_{(i)})$$
$$rloss_{(i+1)} = rloss_{(i)} + \alpha loss * (lossout_{(i)} - rloss_{(i)})$$

Where $rwin_{(i+1)}$ and $rloss_{(i+1)}$ are the estimated probabilities of the win or loss on trial *i+1* being associated with shape A. These probability estimates are the sum of the probability estimates on the previous trial, $rwin_{(i)}$ and $rloss_{(i)}$, and the prediction error on the previous trial weighted by the win or loss learning rate, αwin or $\alpha loss$, which were allowed to vary between 0 and 1. The prediction error on trial *i* is the difference between the estimated probability of the win or loss

occurring with shape A on trial *i*, and the actual outcome $winout_{(i)}$ or $lossout_{(i)}$ on trial *i*. The actual outcomes were coded as 1 if the outcome was associated with shape A, and as 0 if the outcome was associated with shape B.

A Softmax function was used to transform the trial-wise probability estimates into choice probabilities:

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

Where *P(choice = A)$_{(i)}$* represents the probability of the participant choosing shape A on trial *i*. The inverse decision temperature β captures choice stochasticty. High β values indicate that participants were largely influenced by the probability estimates, whereas lower β values suggest more random choice behaviour. The tendency parameter *t* accounts for a potential tendency of choosing one shape over the other.

Model 2. Alternatively, participants might over- or underestimate the probability of wins being associated with shape A vs. losses being associated with shape A. Therefore, the Softmax function in model 2 contains two separate inverse temperatures for wins and losses that independently scale the estimated probability of wins and losses.

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

The probability estimates $rwin_{(i+1)}$ and $rloss_{(i+1)}$ are calculated as in model 1.

Model 3. While model 1 and 2 assume that participants maintain separate probability estimates for wins and losses, participants might simply learn the overall value of the shapes. In model 3, the overall value of shape A on trial *i+1*, $v^A{}_{(i+1)}$, is calculated by

$$v^{A}{}_{(i+1)} = v^{A}_{(i)} + a * (out_{(i)} - v^{A}_{(i)})$$

In this model, the overall value of shape A on trial *i* is updated by the product of a single learning rate parameter a and the prediction error $(out_{(i)} - v^{A}_{(i)})$. $out_{(i)}$ represents the overall outcome and is coded as 1, -1 or 0, depending on whether shape A was associated with a win, a loss, both or neither outcome.

Model 4. Model 4 is identical to model 1 but does not include a tendency parameter in the Softmax function:

$$P(choice = A)_{(i)} = \frac{1}{1 + exp^{(-\beta * (rwin_{(i)} - rloss_{(i)}))}}$$

Model 5. In model 5, the estimated probabilities of the win or loss being associated with shape A are updated using the same learning rate parameter α:

$$rwin_{(i+1)} = rwin_{(i)} + \alpha * (winout_{(i)} - rwin_{(i)})$$
$$rloss_{(i+1)} = rloss_{(i)} + \alpha * (lossout_{(i)} - rloss_{(i)})$$

The estimated probabilities are transformed into choice probabilities using the same Softmax function as in model 2, i.e. separate inverse temperature parameters for win and loss outcomes.

Model 6. Model 6 is a modification of model 2, in which the estimated probabilities of the wins and losses being associated with shape A are centred around zero:

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

The two inverse temperature parameters βwin and $\beta loss$ scale the difference of the estimated probabilities to 0.5 which can capture if individuals are more sensitive to one outcome than to the other.

Model comparison showed that model 1 provided the best model fit to participants' choices (Figure 4.3). The model included 4 parameters, the win learning rate α_{win}, loss learning rate α_{loss}, inverse temperature β and tendency parameter *t*. The model was fitted separately to each of six task blocks for each session. This yielded 12 estimates for each parameter for each participant which were entered into the statistical analysis.

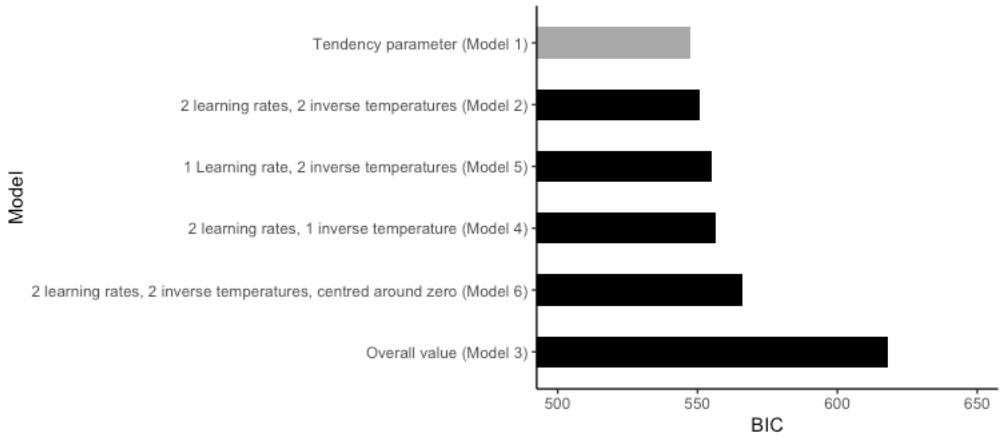

Figure 4.3. Model comparison. The six different models were fitted to each participant and each task block. The winning model was model 1 (tendency parameter).

4.2.8 Parameter estimation

Learning rate parameters were allowed to vary between 0 and 1, the inverse temperature between 0.1 and 100, and the tendency parameter between -1 and 1. The range for each parameter was divided into 30 equally spaced discrete values. For a model with *n* parameters, an *n*-dimensional matrix was created with 30^n entries. For each of the 30^n possible parameter value combinations, the model was used to simulate on a trial-by-trial basis the likelihood given this particular combination of parameter values. By multiplying the trial-by-trial likelihood of the participants' choices given

the combination of parameter values, the posterior probability of the participant's behaviour given this combination of parameter values was calculated. This procedure was repeated for each possible combination of parameter values. The best estimate for a parameter was defined as the expected value of this parameter's marginalised probability distribution. Each parameter's marginalised probability distribution was created by summing across all other dimensions. The expected value was extracted by summing up the discrete parameter values multiplied by their marginalised probabilities.

For statistical analysis, parameter estimates were transformed to approximate a normal distribution. An inverse logit transformation was applied to the learning rate estimates projecting them onto the infinite real line. Inverse temperature estimates were log-transformed. To facilitate interpretation, figures display the non-transformed learning rate and inverse temperature estimates.

To test whether our procedure of parameter estimation can reliably reproduce parameter values from the winning model, parameter recovery was performed for model 1. First, choices were simulated from model 1 using 1000 combinations of parameter values. For each of the four parameters, 10 equally spaced values covering the entire parameter range considered in the estimation procedure were used for simulating choices. Model 1 was than fitted to the simulated choices and the resulting parameter estimates were compared to the actual parameter estimates used for simulating the choices. Overall, the results indicate that the estimated parameters were reasonably close to the actual parameter values used for simulation. For the inverse temperature, there was a ceiling effect at around 40, indicating that parameter values above 40 do not cause changes in behaviour. More details are included in Figure 4.4.

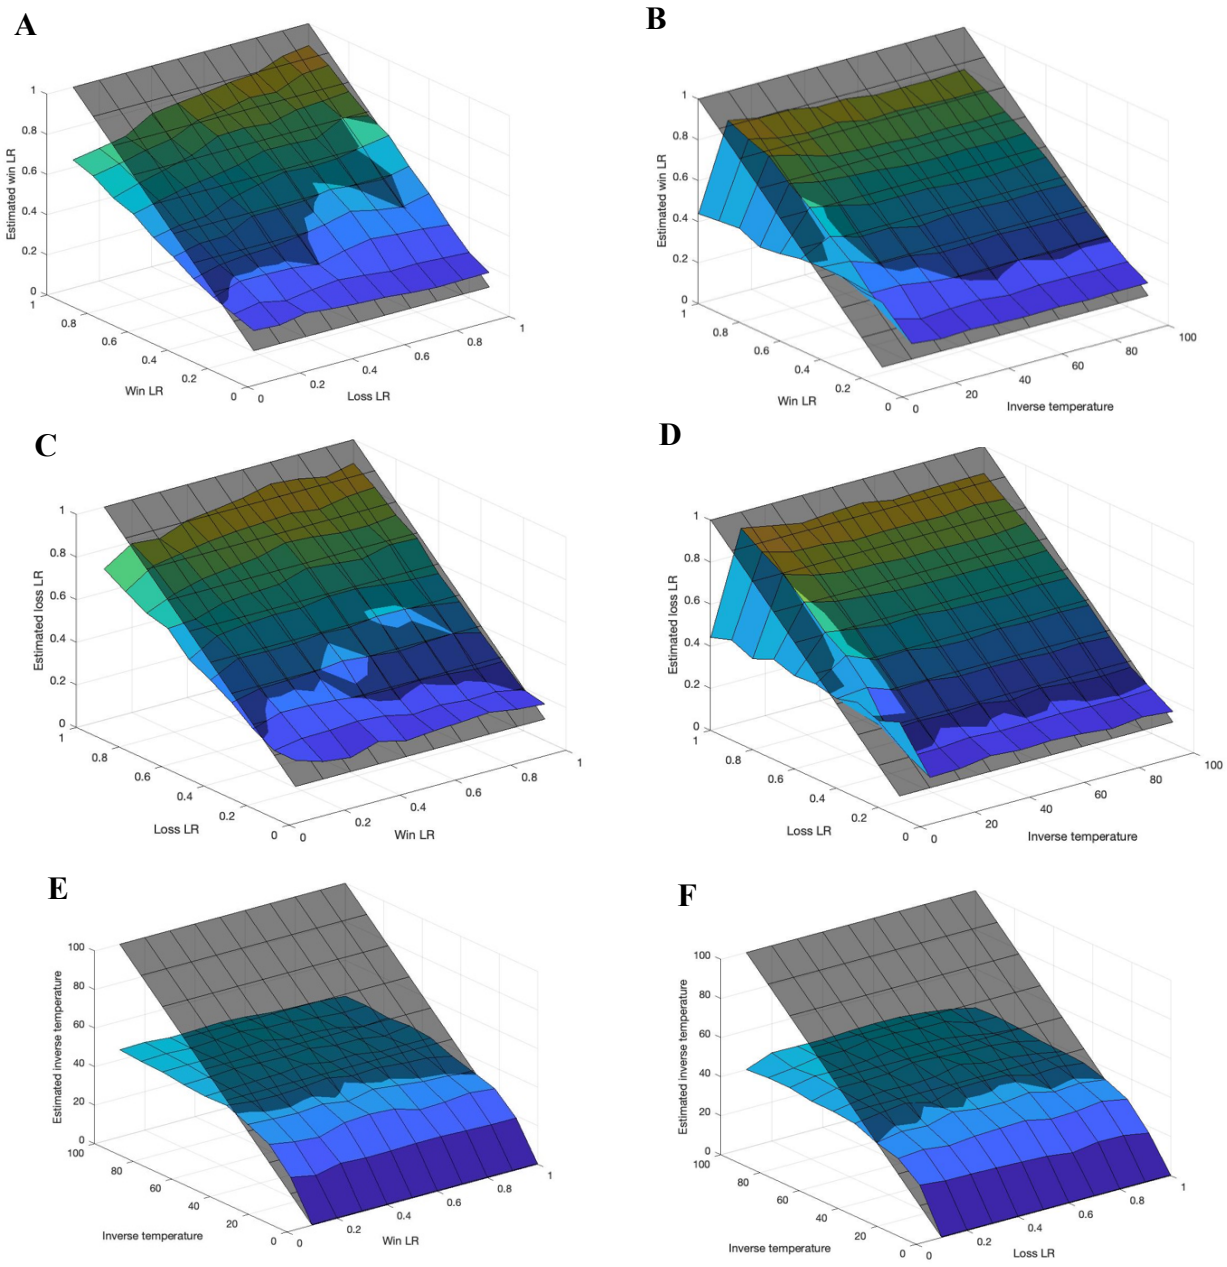

Figure 4.4. Parameter recovery for model 1. Recovery of win learning rate depending on loss learning rate (A) or inverse temperature (B). Recovery of the loss learning rate depending on the win learning rate (C) or inverse temperature (D). Recovery of the inverse temperature depending on the win learning rate (E) or loss learning rate (F). The learning rate for one outcome tends to be underestimated if the learning rate for the other outcome is low (A and C). Learning rates cannot reliably be estimated if the inverse temperature is very low (B and D). There was a ceiling effect for the inverse temperature parameter suggesting that values above a certain threshold (around 40) do not cause different behaviour in the task (E and F).

4.2.9 Additional outcome measures

As a non-computational measure of affective biases, the proportion of "win-driven choices" was calculated. This measure was based on choices that participants made after a trial in which the win and loss occurred with the same shape. These trials are of particular interest because they can differentiate between two opposite strategies: If participants are primarily trying to maximise the number of wins they receive, they should choose the shape that was associated with both outcomes on the previous trial. If participants are mainly trying to avoid losses, they should pick the shape that was associated with neither of the outcomes. The proportion of "win-driven choices" is therefore the number of trials in which participants chose the shape that was associated with both outcomes, divided by the total number of trials in which the win and loss were associated with the same shape.

4.2.10 Statistical analysis

All analyses were performed in RStudio (Version 1.4.1717, R 4.1.1). Choice behaviour in the IBLT was analysed in repeated-measures analysis of variance (ANOVA), with the main factor of interest being tDCS Condition (real vs. sham tDCS). Since tDCS was applied after the first both-volatile block, both-volatile blocks were analysed separately from wins-volatile and losses-volatile blocks. The main outcome measures of interest were learning rates and the proportion of win-driven choices. Although we did not expect tDCS to affect the inverse temperature, exploratory analyses on the inverse temperature estimates were run. P-values below 0.05 are considered significant, whereas p-values between 0.05 and 0.10 are interpreted as trends.

To test our first hypothesis that real bifrontal compared to sham tDCS applied during task performance would selectively increase win learning rates and/or decrease loss learning rates, our first analysis focused on the "online tDCS" group. The primary measure of interest was the learning rate during the wins-volatile and losses-volatile blocks. The learning rate estimates were

entered into a repeated-measures ANOVA which included the within-subject factors tDCS Condition (real vs. sham tDCS), Valence (win vs. loss learning rate), Volatility (wins-volatile vs. losses-volatile), and Time (during vs. after tDCS). The effect of interest was the interaction between tDCS Condition and Valence which was followed up with planned comparisons contrasting real and sham tDCS separately for win and loss learning rates. All analyses included the Block Order (wins-volatile first vs. losses-volatile first) and tDCS Order (real first vs. sham first) as between-subject factors of no interest. All analyses of learning rates included the win and loss learning rates during the first both-volatile block in the first session as a covariate to control for potential differences in baseline learning rates.

To analyse whether online bifrontal compared to sham tDCS increased the proportion of win-driven choices, the proportion of win-driven choices in wins-volatile and losses-volatile blocks was entered into a repeated measures ANOVA including the factors tDCS Condition, Volatility and Time. We tested for a main effect of tDCS Condition and potential interactions including the factor tDCS Condition. The inverse temperature estimates were analysed in the same way but were not hypothesised to be influenced by tDCS.

Learning rates and proportion of win-driven choices in the both-volatile blocks were analysed in a repeated-measures ANOVAs including the same factors apart from Volatility. Since the first both-volatile block was performed before the application of tDCS, and the second both-volatile block was performed 15mins after the stimulation had ended, this analysis was exploratory.

To test our second hypothesis, that the effect of tDCS is cognitive-state dependent, data of the "online tDCS" group and "offline tDCS" group were entered into a mixed ANOVA including the between-subject factor Cognitive State (online vs. offline tDCS). The within-subject factors were the same as in the previous ANOVA, i.e. tDCS Condition, Valence, Volatility and Time. The effect of interest was the three-way interaction between Cognitive State, tDCS Condition and Valence. Win-driven choices were analysed in a mixed ANOVA including the same factors apart

from Valence. We tested for an interaction effect between Cognitive State and tDCS Condition. Inverse temperature estimates were analysed in an identical ANOVA.

To test our third hypothesis, that the effect of tDCS would be anatomically specific, learning rate estimates for the "online tDCS" and "M1" groups were jointly analysed in a mixed ANOVA including the between-subject factor Target Region (bifrontal vs. M1 tDCS). As in the previous ANOVAs, tDCS Condition, Valence, Volatility and Time were included as within-subject factors. The three-way interaction between Target Region, tDCS Condition and Valence was of particular interest. The effect of tDCS on the proportion of win-driven choices was analysed in a mixed ANOVA including the same factors excluding Valence. An identical ANOVA was run on the inverse temperature estimates.

4.3 Results

4.3.1 Online bifrontal tDCS increases win learning

To test our main hypothesis that real bifrontal compared to sham tDCS would increase learning from win outcomes and/or decrease learning from loss outcomes, we conducted a repeated measures ANOVA including the factors tDCS Condition, Valence, Volatility and Time. In line with previous studies, we observed a significant interaction between Valence and Volatility ($F(1,18) = 94.91$, $p < .001$, $\eta^2_G = .205$), indicating that participants adjusted their learning rates to the volatility context (Figure 4.5). Paired t-tests confirmed that participants showed a higher learning rate for losses than for wins in the losses-volatile block ($t(19) = 3.401$, $p = .002$) and a higher learning rate for wins than for losses in wins-volatile blocks ($t(19) = -7.52$, $p < .001$), indicating that participants preferentially learn from the most informative (volatile) outcome.

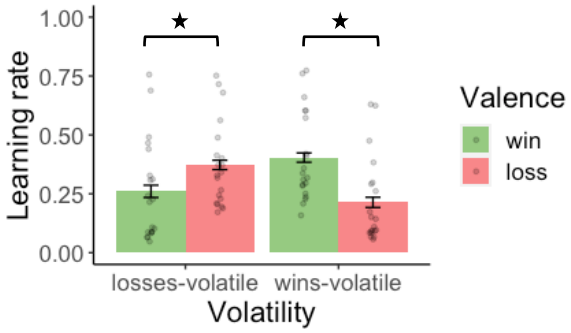

Figure 4.5. Effect of the volatility conditions on learning rates. As expected, participants adjusted their learning rates to the volatility context. In 'losses-volatile' blocks, the learning rate was higher for loses than for wins, and vice versa in the 'wins-volatile' blocks. Error bars indicate SEM. Asterisks indicate a significant effect ($p < .05$).

There was a significant interaction effect between tDCS Condition and Outcome (Figure 4.6A), indicating that 'online' bifrontal tDCS modulated learning rates in a valence-specific way ($F(1,18) = 4.88$, $p = .041$, $\eta^2_G = .007$). Planned comparisons showed that tDCS increased learning from wins ($t(19) = 2.109$, $p = .048$, $d = 0.358$) but not from losses ($t(19) = 0.35$, $p = .72$, $d = 0.067$). Moreover, there was a significant interaction effect between tDCS Condition and Volatility ($F(1,18) = 9.09$, $p = 0.010$, $\eta^2_G = .016$). TDCS increased learning in losses-volatile blocks ($t(19) = 2.46$, $p = .023$), but not in wins-volatile blocks ($t(19) = -0.12$, $p = .90$). The interaction effects of tDCS Condition were not modulated by Time, indicating that the effects did not significantly decrease after the stimulation had finished (tDCS Condition x Outcome x Time: $F(1,18) = 0.55$, $p = .464$, $\eta^2_G = < .001$; tDCS Condition x Volatility x Time: $F(1,18) = 1.25$, $p = .276$, $\eta^2_G = .001$). Although the tDCS Condition x Outcome x Volatility interaction effect was not significant ($F(1,18) = 2.40$, $p = .138$, $\eta^2_G = .002$), the two observed 2-way interaction effects suggest that bifrontal compared to sham tDCS only increased learning from wins in blocks where losses were volatile. To verify this, an ANOVA including only the win learning rate showed that the effect of tDCS on the win learning rate depended on Volatility (tDCS Condition x Volatility: $F(1,18)$, $p = .008$, $\eta^2_G = .031$). Paired t-tests confirmed that tDCS increased win learning rates in the losses-volatile ($t(19) = 2.88$, $p = .009$, $d = .517$) but not in the wins-volatile block ($t(19) = 0.24$, $p = .806$, $d = 0.05$). There was no effect of tDCS on the inverse temperature (main effect of tDCS Condition and all interactions including tDCS Condition $p > .13$).

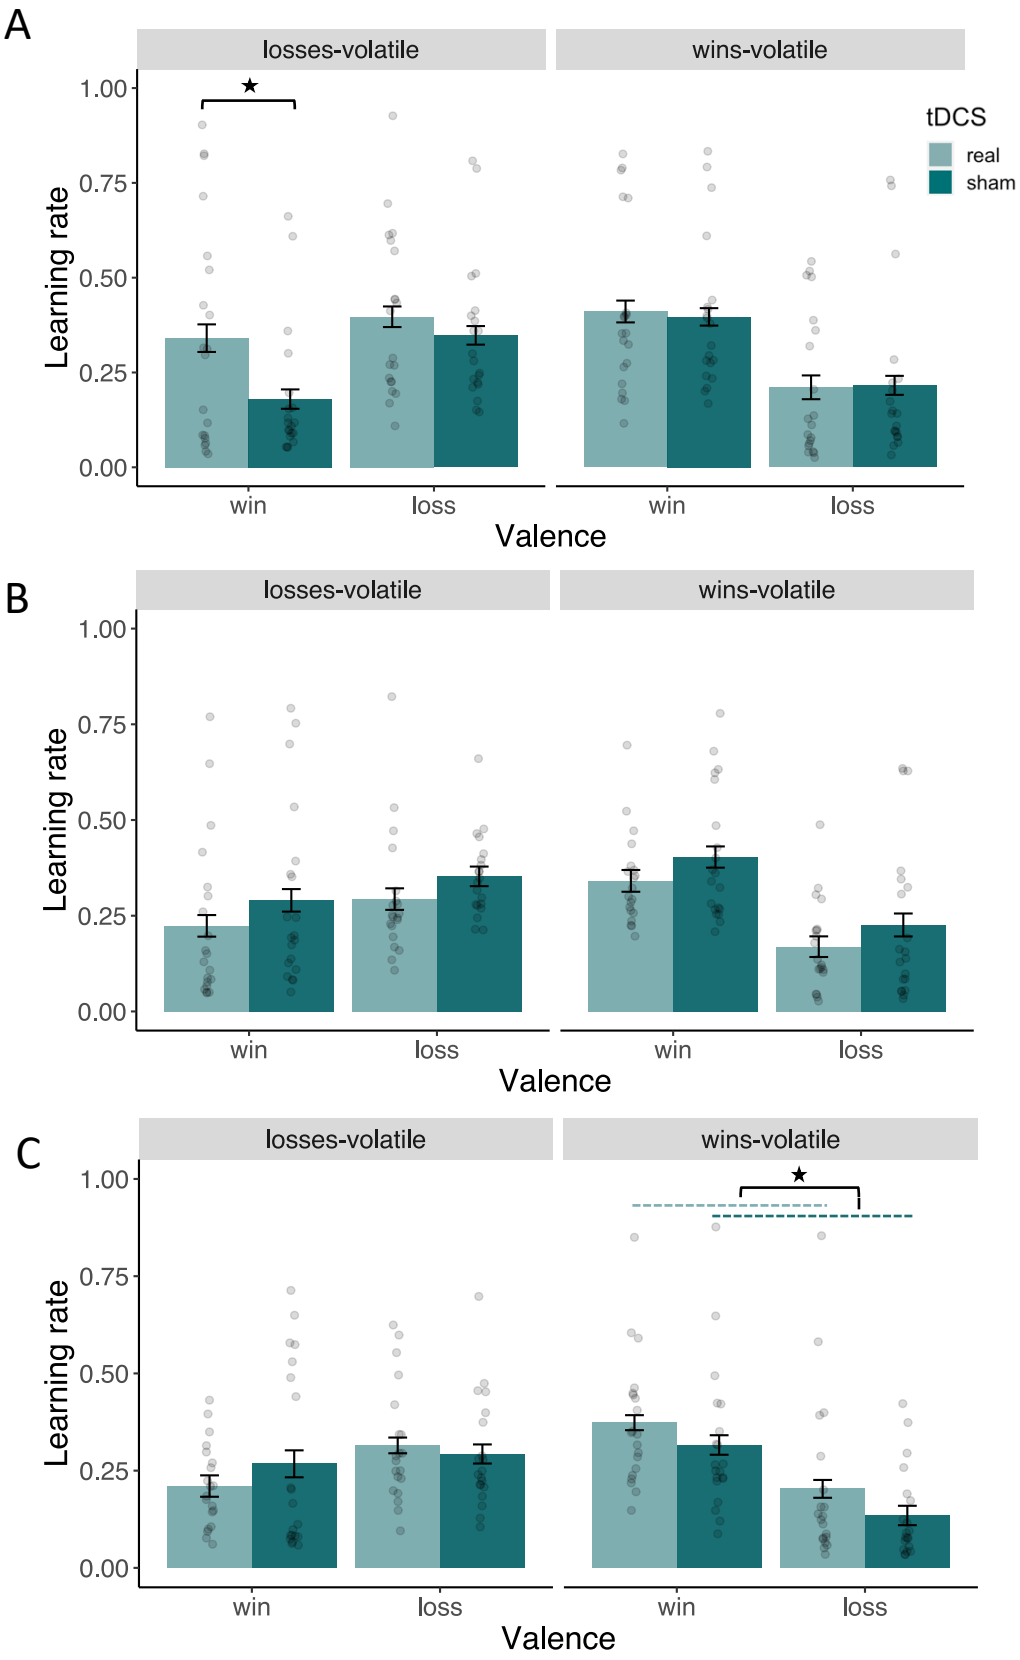

Figure 4.6. Effects of tDCS in the (A) "online bifrontal", (B) "offline bifrontal" and (C) "M1" conditions. (A) Online bifrontal tDCS increased win learning rates in the losses-volatile condition. (B) Offline bifrontal tDCS decreased learning rates in general (across all conditions). (C) TDCS applied to the motor cortex increased learning rates when wins were volatile.

To test our second hypothesis that tDCS would increase the proportion of win-driven choices, we conducted an ANOVA including tDCS Condition, Volatility and Time. There was a significant main effect of tDCS Condition ($F(1,18) = 5.10$, $p = .036$, $\eta^2_G = .014$), and interaction between tDCS Condition and Volatility ($F(1,18) = 7.86$, $p = .011$, $\eta^2_G = .029$), indicating that the effect of tDCS depended on the volatility condition (Figure 4.7A). Post-hoc paired t-tests confirmed that tDCS increased the proportion of win-driven choices in the losses-volatile ($t(19) = 3.31$, $p = .003$, $d = 0.476$) but not in the wins-volatile block ($t(19) = -0.70$, $p = .489$, $d = -0.143$).

To test if tDCS affected task performance, we analysed the amount of money won per block in an ANOVA including tDCS Condition, Volatility and Time. Real compared to sham tDCS significantly reduced the amount of winnings per block ($F(1,18) = 5.90$, $p = .025$, $\eta^2_G = .022$)(Figure 4.7B).

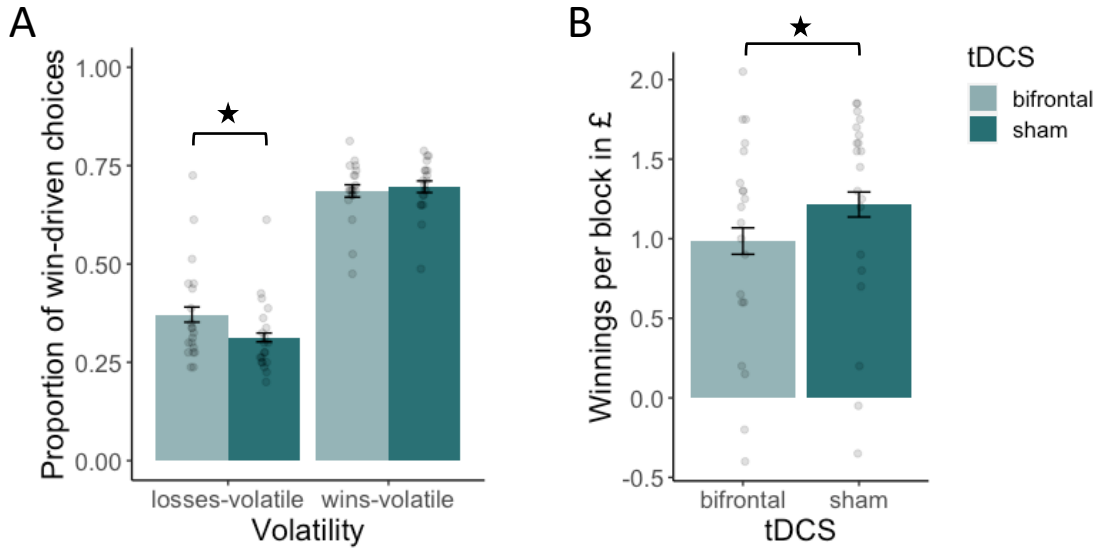

Figure 4.7. Effects of online tDCS on non-computational measures. (A) Real compared to sham tDCS increased the proportion of win-driven choices in blocks in which losses are volatile. (B) Real compared to sham tDCS decreased the amount of winnings.

Regarding both-volatile blocks, an effect of tDCS would be indicated by an interaction between tDCS Condition and Time, since the first both-volatile block was performed before, and the second

after tDCS. There were no interaction effects between tDCS Condition and Time on any computational or non-computational measure in both-volatile blocks.

4.3.2 The effect of online bifrontal tDCS is cognitive-state dependent

To test whether the effect of tDCS on learning from positive outcomes is cognitive-state dependent, i.e. specific to the time point of stimulation ('online' vs. 'offline'), we conducted an ANOVA on the combined dataset of 'online' and 'offline' tDCS including the between-subject factor Cognitive State, and the within-subject factors tDCS Condition, Valence, Volatility and Time as in the analyses above. There was an interaction effect between Cognitive State and tDCS Condition ($F(1,36) = 6.26$, $p = .016$, $\eta^2_G = .018$). Post-hoc t-tests indicated that 'offline' compared to sham tDCS led to decreased learning rates in general ($t(19) = -2.17$, $p = .042$) whereas 'online' tDCS had no such effect (Figure 4.6B)($t(19) = 1.37$, $p = .183$).

Moreover, there was a significant interaction effect between Cognitive State, tDCS Condition and Valence ($F(1,46) = 5.10$, $p = .036$, $\eta^2_G = .002$) and a significant interaction effect between Cognitive State, tDCS Condition and Volatility ($F(1,36) = 8.07$, $p = .007$, $\eta^2_G = .006$). Follow-up ANOVAs indicated that, in contrast to the 'online' bifrontal sample, the interaction effects including tDCS Condition were not significant in the 'offline' bifrontal sample (tDCS Condition x Outcome: $F(1,18) = 0.63$, $p = .435$, $\eta^2_G = .0005$; tDCS Condition x Volatility: $F(1,18) = 0.97$, $p = .335$, $\eta^2_G = .001$). This suggests that the effects of bifrontal tDCS on learning are specific to the cognitive state, since they were present in participants who received 'online' but not those who received 'offline' tDCS. Taken together, 'online' tDCS selectively increased learning from wins when losses where volatile, whereas 'offline' tDCS decreased learning rates in general.

To test our hypothesis that the effect of tDCS on win-driven choices would also be cognitive-state dependent, we conducted an ANOVA on the combined dataset including the factors Cognitive State, tDCS Condition, Volatility and Time. There was a significant interaction effect between

Cognitive State, tDCS Condition and Volatility ($F(1,36) = 8.37$, $p = .006$, $\eta^2_G = .017$) suggesting that the effect of tDCS on the proportion of win-driven choices differed depending on whether stimulation was delivered before or during task performance. In contrast to online tDCS, offline tDCS had no significant effect on win-driven choices, although there was a trend towards an increase in win-driven choices which did not depend on the volatility condition (main effect of tDCS Condition: $F(1,18) = 3.52$, $p = .076$, $\eta^2_G = .008$), tDCS Condition x Volatility interaction: $F(1,18) = 1.85$, $p = .190$, $\eta^2_G = .008$). This confirms our hypothesis that the tDCS effect on win-driven choices is cognitive-state dependent. Offline bifrontal tDCS did not affect the amount of winnings (all $p > .05$).

4.3.3 The effect of 'online' bifrontal tDCS is anatomically specific

To test our hypothesis that the effect of 'online' tDCS would be anatomically specific, we conducted an ANOVA combining the 'online' bifrontal and motor cortex samples. There was a significant interaction effect between Target Region, tDCS Condition and Outcome ($F(1,36) = 4.29$, $p = .045$, $\eta^2_G = .005$) and a significant interaction effect between Target Region, tDCS Condition and Volatility ($F(1,36) = 13.57$, $p < .001$, $\eta^2_G = .014$), suggesting that the interaction between tDCS Condition and Outcome or Volatility was modulated by the target region. A follow-up ANOVA indicated that the interaction effect of tDCS Condition and Outcome was not significant in the motor cortex sample (tDCS Condition x Outcome: $F(1,18) = 1.01$, $p = .326$, $\eta^2_G = .003$). However, the interaction effect between tDCS Condition and Volatility was significant (Figure 4.6C)($F(1,18) = 5.60$, $p = .029$, $\eta^2_G = .013$). Post-hoc t-tests indicate that M1 tDCS increased learning when wins were volatile ($t(19) = 2.40$, $p = .026$) but not when losses were volatile ($t(19) = -0.37$, $p = .710$). In contrast to 'online' bifrontal tDCS, this effect was not valence-specific (i.e. it did not have a specific effect on learning from wins).

An ANOVA was conducted on the proportion of win-driven choices to test whether the increase in the proportion of win-driven choices induced by 'online' bifrontal tDCS was specific to the

bifrontal tDCS. There was a significant interaction between Target Region, tDCS Condition and Volatility ($F(1,36) = 7.11$, $p = .011$, $\eta^2_G = .018$) , suggesting that the effect of tDCS depended on the target region. A post-hoc ANOVA in the motor cortex group confirmed that tDCS applied to M1 did not have an effect on win-driven choices (main effect of tDCS: $F(1,18) = 0.04$, $p = .831$, $\eta^2_G < .001$); all interactions involving tDCS Condition $p > .24$). TDCS applied to the motor cortex did not affect the amount of winnings (all $p > .05$).

4.3.4 Replication of the valence-specific effect of 'online' bifrontal tDCS

To test whether the increase in win learning rate and win-driven choices induced by 'online' bifrontal tDCS was replicable in an independent sample, we recruited another sample of the same size (n = 20) that underwent the same experimental procedure as the original sample. We started by replicating the same analyses that were conducted in the original dataset. In contrast to the original sample, in this sample an ANOVA conducted on the learning rates including the factors tDCS Condition, Valence, Volatility and Time did not find a significant tDCS Condition x Valence interaction ($F(1,18) = 1.30$, $p = .268$, $\eta^2_G = .002$) or tDCS x Volatility interaction ($F(1,18) = 0.01$, $p = .929$, $\eta^2_G < .001$). Planned comparisons showed that real vs. sham tDCS had no effect on win ($t(19) = 1.02$, $p = .32$) or loss learning rates ($t(19) = 0.34$, $p = .73$). Since the increase in win learning rate in the original sample was only found in the losses-volatile condition, we ran a t-test contrasting real vs. sham tDCS on the win learning rate in the losses-volatile condition which was non-significant ($t(19) = 0.99$, $p = .33$). In addition, we tested whether tDCS had an effect on win or loss learning rates during the stimulation period (block 2 and 3). TDCS had a marginally significant effect on the win learning rate *during* ($t(19) = 2.08$, $p = .050$, $d = 0.387$) but not after the stimulation period ($t(19) = 0.17$, $p = .86$)(Figure 4.8A). This effect did not depend on the volatility condition (ANOVA including the factors tDCS Condition and Volatility on the win learning rate during stimulation, tDCS Condition x Volatility interaction: $F(1,19) = 0.27$, $p = .61$). No significant effect on the loss learning rates was observed (during tDCS: $t(19) = 1.1$, $p = .28$; after tDCS: $t(19) = -0.62$, $p = .52$).

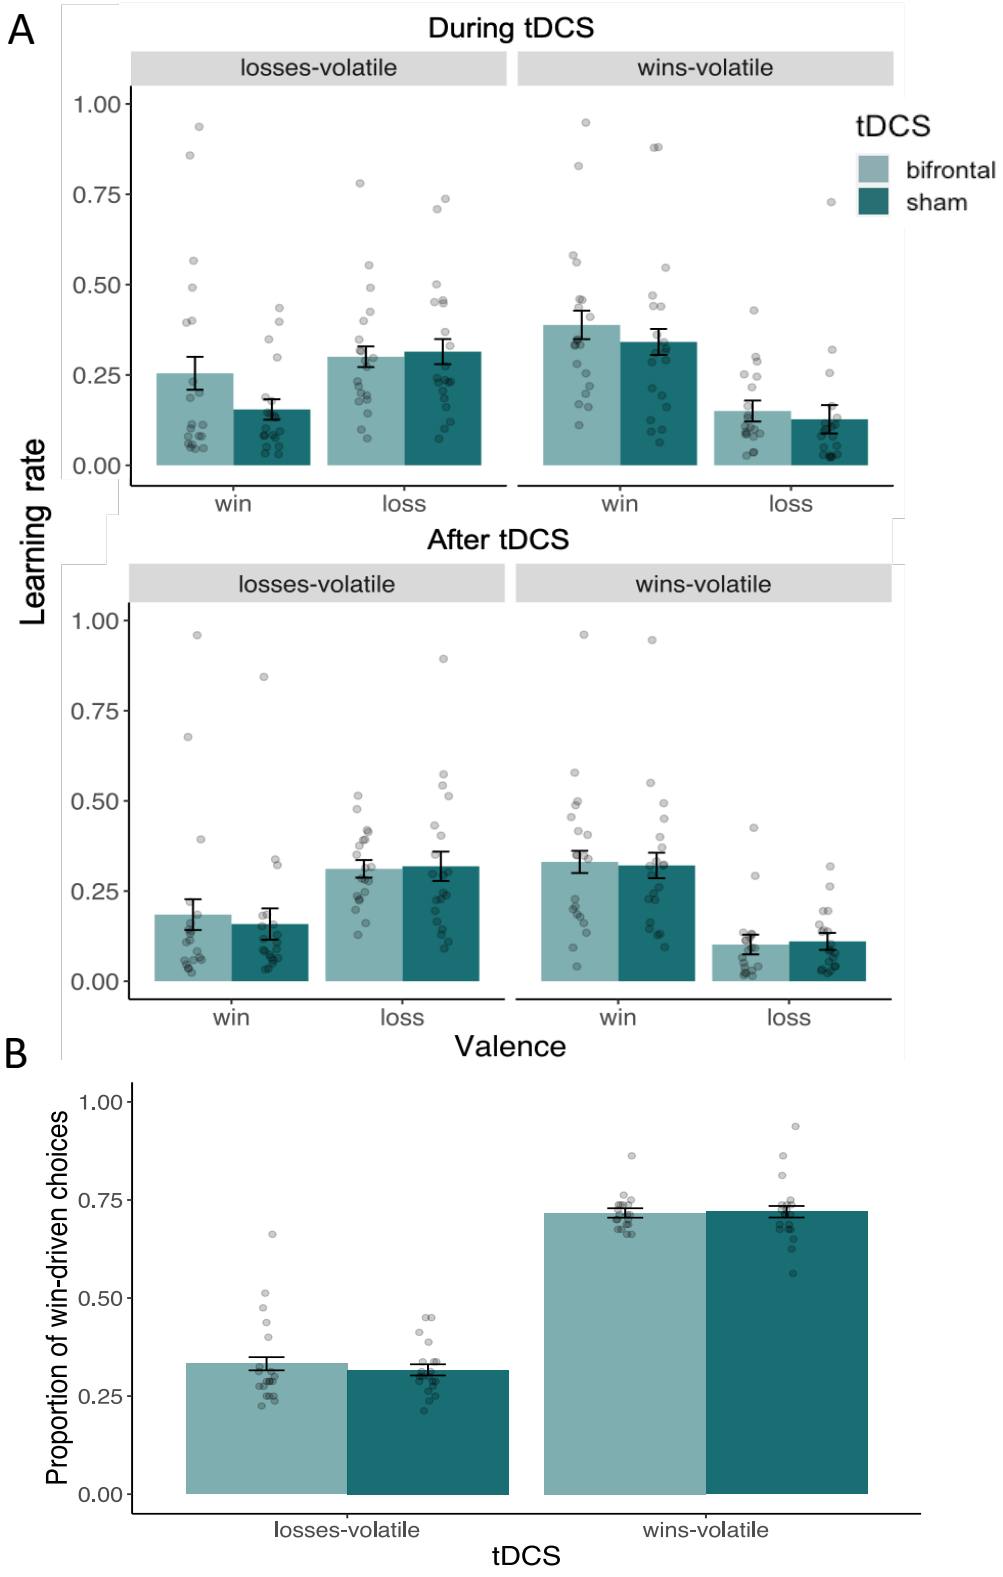

Figure 4.8. Replication of the effects of 'online' bifrontal tDCS in an independent sample. (A) In the replication sample, tDCS led to a marginal increase in win learning rates only during the stimulation period (p = .05, upper panel). (B) TDCS did not increase the proportion of win-driven choices in the replication sample.

To test whether the effect of 'online' bifrontal tDCS on win-driven choices would replicate in the independent sample, an ANOVA including tDCS Condition, Volatility and Time was conducted with the proportion of win-driven choices as dependent variable. There was no significant main effect of tDCS Condition ($F(1,18) = 0.26$, $p = .615$, $\eta^2_G = .001$) or interaction effect involving tDCS Condition (all $p > .11$)(Figure 4.8B).

Taken together, the effects of 'online' bifrontal tDCS observed in the original sample partially replicated in an independent sample. In the original sample, 'online' bifrontal tDCS selectively increased win learning rates for the duration of the experiment. In the replication sample, however, there was only evidence for an increase in win learning rates during the stimulation period, indicating that the effect was less long-lasting. In the original sample, the increased focus on win outcomes also manifested on non-computational measures as an increase in the proportion of win-driven choices. In contrast, in the replication sample there was no evidence for an effect on win-driven choices.

4.4 Discussion

In this study, we investigated whether bifrontal tDCS would increase processing of positive vs. negative information in the Information Bias Learning Task. As hypothesised, we found that bifrontal tDCS specifically increased win learning rates, which lasted for at least 15 minutes after cessation of the stimulation. In a replication study, a similar effect was observed, albeit weaker and less long-lasting. The valence-specific effect of tDCS was specific to the target region (bilateral DLPFC not M1) and was dependent on the cognitive state of participants at the timepoint of stimulation (i.e. tDCS applied during vs. before task performance).

The reversal of negative biases is thought to be one key mechanism underlying effective antidepressant treatments (Godlewska et al., 2016; Godlewska et al., 2012; Harmer et al., 2009). We found that in healthy individuals online bifrontal tDCS induced a selective increase in reward

learning. In line with this, online bifrontal tDCS increased a non-computational measure of positive bias. These findings suggests that online bifrontal tDCS can increase the processing of positive information, an effect of potential interest to counter-act negative biases in depression. While online bifrontal tDCS had a negative effect on task performance in healthy individuals (reduced winnings), the increase in reward learning might be beneficial in individuals suffering from depression who might show a bias towards punishment learning (Aylward et al., 2019; Eshel & Roiser, 2010; Pike & Robinson, 2022).

The valence-specific effect of bifrontal tDCS is in line with findings from previous studies. Anodal tDCS applied to the left DLPFC has previously been found to improve cognitive control for positive stimuli (Vanderhasselt et al., 2013) and recognition of positive emotions (Nitsche et al., 2012), and reduce attention to negative stimuli (Brunoni et al., 2014; Chen et al., 2017; Ironside et al., 2019; Ironside et al., 2016). Bifrontal tDCS might shift the balance away from the right hemisphere associated with processing of negative stimuli, towards the left hemisphere associated with the processing of positive stimuli (Hecht, 2010; Herrington et al., 2005; Wyczesany, Capotosto, Zappasodi, & Prete, 2018).

The DLPFC is part of a network involved in reinforcement learning (Haber & Knutson, 2010). Increasing evidence suggests that the DLPFC is involved in adjusting learning rates to volatility (Farashahi et al., 2019; Massi et al., 2018). The observed effect might also be caused by excitability changes in the dorsal anterior cingulate, an anatomically connected area which has been shown to be activated in response to volatility (Behrens et al., 2007). Electric field modelling suggests that the highest field strength might actually be generated in between rather than underneath the electrodes (Karabanov et al., 2019), i.e. in the medial prefrontal cortex. The medial prefrontal cortex has been associated with different reinforcement learning processes, including action-outcome associations and belief updating (Alexander & Brown, 2011; Rushworth, 2008). Future studies could combine tDCS with neuroimaging to investigate how the behavioural effect of tDCS relates to physiological changes in these brain regions.

The effect of bifrontal tDCS on reward learning rates depended on the volatility context: it was only observed when losses were volatile and therefore more informative than wins. One possible explanation for this might be a ceiling effect in blocks where wins were volatile, since the learning rates from wins were already very high in these blocks during sham tDCS. The effect of tDCS in the replication sample, however, did not depend on the volatility condition.

The valence-specific effect of online bifrontal tDCS was cognitive-state dependent. In contrast to online bifrontal tDCS, offline bifrontal tDCS led to decreased learning rates in general. This is consistent with the known brain-state dependency of non-invasive brain stimulation effects (Li et al., 2019; Silvanto, Muggleton, & Walsh, 2008). Although we do not fully understand why offline tDCS caused a decrease in learning rates, differential effects of online and offline tDCS have been observed previously (Li et al., 2019; Stagg et al., 2011) and are hypothesised to depend on metaplastic mechanism (Müller-Dahlhaus & Ziemann, 2015). While tDCS in depression trials is typically applied at rest, the cognitive-state dependency observed in our study suggests that stimulating the brain during task performance could potentially have greater functional impact by modulating cognitive factors (increased reward learning) believed to causally maintain depression (negative bias). TDCS has been shown to increase plasticity and learning effects (Clarke et al., 2014; Fritsch et al., 2010; O'Shea et al., 2017; Reis et al., 2009). Applying tDCS during the performance of a reinforcement learning task designed to induce a positive bias might specifically strengthen the neural circuits engaged in reward processing and might counter-act negative biases.

The effect of bifrontal tDCS was specific to the target region. tDCS applied to the bilateral motor cortex did not induce any valence-specific changes. Unexpectedly, we found that motor cortex tDCS increased learning in wins-volatile blocks in this dataset. Since tDCS stimulates large areas of the brain, we speculate that this effect was caused by activity changes in neighbouring areas such as the anterior cingulate (Behrens et al., 2007).

As a limitation of the study, it should be noted that we did not correct for multiple comparisons. In the online bifrontal sample, only one hypothesis (increased reward vs. punishment learning) was tested. All other analyses in this sample were control analyses. To test for the specificity of the effect, we recruited two additional independent samples that were tested in two control conditions (cognitive state and control region). Adding these control conditions increases the overall likelihood of obtaining false-positive results, but it does not specifically increase the likelihood that the effect observed in the online bifrontal sample is a false-positive finding. However, the effects observed in the two control conditions should be interpreted with caution since they were unexpected. To test whether the effect observed in the online bifrontal sample might be a false-positive finding, we conducted a replication study. The effect of online bifrontal tDCS did not fully replicate in an independent sample. The interaction between tDCS Condition and Valence did not reach significance, and the increase of win learning rates was only present during the stimulation period. However, it should be kept in mind that only a single session of tDCS was applied in this study. The effects of tDCS are likely to accumulate if applied throughout multiple testing sessions. Another limitation is that we did not test for transfer effects. Hence, it is unclear whether the increase in reward learning would generalise to contexts outside of the experimental paradigm.

The observed increase in reward learning might be beneficial in depression treatment if negative biases in depression are caused by increased punishment learning (Aylward et al., 2019; Eshel & Roiser, 2010; Pike & Robinson, 2022). Although this is a current hypothesis, there is also contradicting evidence (Beevers et al., 2013; Kunisato et al., 2012; Mukherjee, Filipowicz, Vo, Satterthwaite, & Kable, 2020). The next chapter therefore investigates alterations in reinforcement learning in individuals with depressive symptoms. Moreover, it is unclear whether the observed effects of tDCS would generalise to clinical populations. It is possible that in the present study online bifrontal tDCS increased a pre-existing optimism bias typically observed in healthy individuals (Korn et al., 2014; Sharot, Korn, & Dolan, 2011). Chapter 6 therefore investigates the effect of online bifrontal tDCS in individuals with depressive symptoms.

4.5 Conclusions

To conclude, we found that online bifrontal tDCS applied during a reinforcement learning paradigm selectively increased reward learning rates in healthy individuals. This finding supports the hypothesis that the reversal of negative biases might be one potential mechanism underlying the antidepressant effect of bifrontal tDCS. To investigate whether an increase in reward learning might normalise information processing in depression, the next chapter investigates alterations in reinforcement learning in individuals with depressive symptoms.

4.6 Supplementary analyses

Descriptive statistics for the STAI state-anxiety scale and PANAS are included in the table below (Table 4.2). To test whether tDCS had an influence on how participants' affect changed over the course of each session, we analysed the scores in repeated-measures ANOVAs including the factors tDCS Condition (real vs. sham tDCS) and Time (Before and after the session). An interaction effect between tDCS Condition and Time would indicate that tDCS had an effect on participants' affect. A positive interaction was only found for the PANAS negative affect score in the offline tDCS group ($F(1.19) = 5.7$, $p = .02$). While negative affect decreased over the course of the session in the sham tDCS condition ($F(1,19) = 11.4$, $p = .003$), no change was observed in the real tDCS condition $F(1,19) = 0.1$, $p = .70$).

Table 4.2. State-anxiety and affect scores before and after each testing session, mean (SD).

	Online tDCS		**Offline tDCS**		**Motor cortex tDCS**		**Online tDCS (replic.)**	
	Before	After	Before	After	Before	After	Before	After
STAI-S								
Real	30.0 (6.1)	30.6 (4.6)	29.5 (7.3)	30.8 (8.5)	29.7 (7.8)	28.9 (6.5)	31.8 (6.9)	30.6 (7.8)
Sham	29.9 (4.6)	30.9 (4.8)	32.2 (8.6)	30.6 (8.0)	28.7 (6.5)	29.3 (7.5)	29.7 (6.4)	28.9 (7.2)
PANAS pos								
Real	27.2 (5.3)	24.4 (6.5)	28.6 (6.3)	24.3 (6.0)	30.5 (7.2)	27.6 (8.7)	28.1 (8.5)	31.3 (7.9)
Sham	26.8 (4.8)	22.4 (6.1)	28.1 (6.4)	26.0 (6.4)	31.3 (7.4)	27.5 (7.6)	30.3 (8.6)	32.5 (7.2)
PANAS neg								
Real	11.6 (2.2)	10.6 (0.8)	12.0 (2.6)	11.9 (3.0)	11.6 (2.0)	11.0 (1.4)	10.8 (1.3)	11.9 (4.5)
Sham	10.9 (1.1)	10.4 (0.5)	12.9 (2.9)	11.5 (2.5)	11.6 (2.1)	11.5 (2.3)	10.9 (1.3)	11.0 (1.8)

STAI-S: State-Trait Anxiety Inventory (State-anxiety scale); PANAS: Positive and Negative Affect Scale (positive / negative affect scale).

5 Reward learning in low mood

In our previous study, we found that bifrontal transcranial direct current stimulation (tDCS) selectively increased reward learning in healthy individuals. Negative biases have been hypothesised to result from increased punishment learning. If depression is indeed associated with increased punishment learning, increasing reward learning might counter-act negative biases and might therefore be beneficial in the treatment of depression. The goal of this study was to investigate how reinforcement learning is altered in depression. We hypothesised that individuals with depressive symptoms show increased punishment learning and have difficulties adjusting their learning rates to the informativeness of positive and negative outcomes. In this study 40 healthy individuals and 43 individuals with low mood performed the Information Bias Learning Task, which manipulates the relative informativeness of positive and negative outcomes. In contrast to our first hypothesis, we did not find evidence for increased punishment learning in individuals with low mood. However, low mood was associated with altered learning rate adjustment. While healthy individuals adjusted win and loss learning rates equally to changes in informativeness, participants with low mood adjusted their loss learning rate less than their win learning rate. Compared to healthy individuals, individuals with low mood adjusted their loss learning rate less, and their win learning rate more to changes in informativeness. While a decrease in loss learning rate adjustment is in line with previous literature, the observed increase in win learning rate adjustment was unexpected. We speculate that individuals with low mood had difficulties tracking negative outcomes and allocated more cognitive resources to positive outcomes as a compensatory strategy. We conclude that rather than increasing reward learning, a novel treatment approach could aim at improving individuals' ability to detect changes in the informativeness of positive and negative outcomes and adjusting their choice behaviour accordingly.

5.1 Introduction

In the previous chapter, we found that bifrontal tDCS applied during the Information Bias Learning Task (IBLT)(Pulcu & Browning, 2017) increased reward learning in healthy participants. Since negative biases in depression are hypothesised to result from a bias towards learning from negative outcomes (Aylward et al., 2019; Eshel & Roiser, 2010; Pike & Robinson, 2022) this effect might be beneficial in depression treatment.

Increased punishment learning has been suggested as one potential mechanism that could lead to negative biases in information processing (Aylward et al., 2019; Eshel & Roiser, 2010; Pike & Robinson, 2022). Empirical evidence for this hypothesis is mixed: while some studies found increased punishment (Aylward et al., 2019) or decreased reward learning rates in depression (Cavanagh, Bismark, Frank, & Allen, 2019), other studies did not observe increased punishment vs. reward learning in depression (Beevers et al., 2013; Kunisato et al., 2012; Mukherjee et al., 2020). However, recent research suggests that individuals with depressive symptoms might have difficulties estimating the informativeness of outcomes and adjusting their behaviour accordingly (Browning et al., 2015; Gagne et al., 2020). Negative biases might develop if individuals systematically overestimate the informativeness of negative events (Pulcu & Browning, 2017).

The aim of this study was to investigate negative biases in the IBLT in individuals with depressive symptoms. We compared choice behaviour between individuals with depressive symptoms and healthy individuals in the Information Bias Learning Task. By manipulating the relative informativeness of positive and negative outcomes, the paradigm allows to assess biases towards learning from positive or negative outcomes depending on their relative informativeness, and the extent to which participants adjust reward and punishment learning rates to changes in informativeness. We hypothesised that (1) individuals with depressive symptoms would show increased punishment vs. reward learning rates compared to healthy individuals. Moreover, we

hypothesised that (2) individuals with depressive symptoms would adjust their learning rates to a smaller extent in response to changes in informativeness compared to healthy individuals.

5.2 Methods

5.2.1 Samples

To investigate the influence of depressive symptoms on reward and punishment learning, two samples were analysed: A sample of healthy individuals and a sample of participants with depressive symptoms. Both samples were recruited for tDCS studies (healthy participants for the study reported in chapter 4, and participants with depressive symptoms for the study reported in chapter 6). All participants took part in two sessions in which they received either real or sham tDCS in counter-balanced order. To compare baseline performance between healthy individuals and individuals with depressive symptoms, only data from the sham tDCS condition were analysed. To avoid confounds with learning effects from task repetition, only individuals who received sham tDCS in their first session were included.

The sample of healthy participants included 40 individuals who received sham tDCS in the first session. Participants were excluded if they had any contraindications to brain stimulation, such as neurological conditions, or if they ever had a diagnosis of a mental disorder. Participants were not pre-screened for symptoms of depression or other mental disorders, therefore this sample will be referred to as “general population”. The sample of participants with depressive symptoms was recruited via advertisement explicitly targeting individuals experiencing “low mood”. Low mood was defined as a score of at least 10 on the Beck Depression Inventory II (BDI)(Beck et al., 1996). Participants were asked to fill out an online screening including the BDI and a tDCS safety questionnaire. Participants were excluded from the study if they had any contraindication to tDCS, such as medication (apart from the contraceptive pill), neurological conditions, or a family history of epilepsy, metal implants inside the brain, or current pregnancy. Participants who scored at least

10 on the BDI were invited to an online interview during which their history of mental health was assessed in the Structured Clinical Interview for the Diagnostic and Statistical Manual of Mental Disorders (DSM-5)(First, 2014) conducted by a trained PhD student. To ensure participants were not suffering from bipolar disorder, participants were excluded from the study if they met the diagnostic criteria for a past manic or hypomanic episode. 43 participants with low mood who received sham tDCS in the first session were included in the analysis. The sample showed a wide range of BDI scores (Figure 5.5). Demographic data for both samples are included in Table 5.1.

Table 5.1. Mean (SD) baseline characteristics for the "general population" and "low mood" sample.

	General population (n = 40)	Low mood (n = 43)
Sociodemographic data		
Female (%)	25 (62%)	32 (74%)
Age in years	25.7 (5.5)	24.4 (4.6)
Clinical measures		
STAI-Trait	35.7 (7.2)	56.9 (8.9)
BDI	4.7 (6.5)	26.8 (9.7)

BDI: Beck Depression Inventory-II, score range = 0-63; STAI-Trait: State-Trait Anxiety Inventory (trait anxiety scale), score range = 20-80.

5.2.2 Experimental procedure

Procedures for the "general population" sample are described in chapter 4. The study on low mood largely followed the same procedures but included a pre-screening session. Participants who expressed an interest in participating in the study were asked to complete an online questionnaire including the BDI. Participants were invited to take part in the study if they scored at least 10 on the BDI and did not have any contraindications to tDCS. Participants were then invited to a video call in which the Structured Clinical Interview for DSM-5 (First, 2014) was administered by a trained PhD student. The main purpose of the interview was to exclude participants who might

suffer from bipolar disorder. Participants who met diagnostic criteria for a current or past manic or hypomanic episode were excluded from the study.

Eligible participants were then invited to two testing sessions in the Oxford Centre for Human Brain Activity (OHBA). They were screened for Covid-19 symptoms before entering the building, were asked to sanitise their hands and were given a clinical mask to wear throughout the entire session. The experimental procedure was the same as in our previous study in healthy participants (chapter 4). In brief, each session started with the completion of questionnaires. Participants completed the Beck Depression Inventory II (BDI)(Beck et al., 1996) and the trait-anxiety scale of the State-Trait Anxiety Inventory (STAI)(Spielberger, 1983) in the first session, and the state-anxiety scale of the STAI and the Positive and Negative Affect Schedule (PANAS)(Watson et al., 1988) in both sessions. Afterwards, the tDCS electrodes were set up, with the anode over the left DLPFC (F3 in the international EEG 10-20 system) and the cathode over the right DLPFC (F4). Since this chapter only includes data from the sham tDCS condition, further details regarding tDCS are omitted in this chapter (but can be found in chapter 6). Once the electrode setup was completed, participants received instructions for performing the Information Bias Learning Task (IBLT)(Pulcu & Browning, 2017) which consisted of six task blocks of 80 trials (see section 4.2.4). Participants were assigned to receive either “online” tDCS or “offline” tDCS. Participants in the online tDCS group received the stimulation during the performance of block 2 and 3 of the IBLT. Participants in the offline tDCS group received the stimulation at rest, between block 1 and 2. In each session, participants received either real or sham tDCS, the order of which was counter-balanced across participants. In the sham tDCS condition, the current was ramped down after 30 seconds. After completion of the task, participants again completed the STAI state anxiety scale and PANAS, and an additional questionnaire on adverse effects of the stimulation and strategies applied during task performance.

5.2.3 Computational modelling

In our first study (chapter 4) models were fitted to each block separately, so that separate parameter estimates were obtained for each of the six blocks per session and participant. For our second study, as described in our preregistration (https://clinicaltrials.gov/ct2/show/NCT03393312), we were planning to apply the same modelling approach (in the following referred to as "block-wise models"). In addition, to complement this standard approach, we also used a second modelling approach in which the inverse temperature was estimated across the whole task ("constant inverse temperature models") rather than separately for each block. The reason for this is that the volatility manipulation in the IBLT has been found to induce changes in learning rates (but not inverse temperature) between the different task blocks (Pulcu & Browning, 2017; Pulcu et al., 2019). However, the block-wise modelling approach also allows the inverse temperature parameters to vary between task blocks. By contrast, in our second modelling approach, inverse temperature parameters were instead estimated across the entire task (depending on the model, there was one inverse temperature parameter for wins and one for losses, or one for both outcomes), with the goal of providing a more robust estimate of this parameter. With this second modelling approach, by keeping the inverse temperature parameters constant across the entire task, changes in choice behaviour between task blocks are attributed to changes in learning rates (rather than changes in inverse temperature). The following sections explain these two modelling approaches in more detail and evaluate model comparison for each of them to identify the models that best fit the observed data.

5.2.3.1 Block-wise models

The six block-wise models investigated for study 2 are the same as in study 1 and are described in section 4.2.7. Each model was fit to each block separately (omitting the first 10 trials of each block), leading to six estimates per parameter per participant and session. Parameter estimates were extracted by calculating the posterior probability over the full parameter space and extracting the

expected value of each parameter's marginalised probability distribution (this procedure is described in section 4.2.8).

Model comparison between the six models was performed for the data collected in study 2 (low mood) to assess whether the model that best describes the data would be the same as in study 1 (general population). The results show that this was the case since model 1, including two learning rates, a tendency parameter and one inverse temperature, also had the lowest BIC in study 2 (Figure 5.1). Parameter recovery for the winning model is included in chapter 4.2.8. All subsequent analyses reported for the "block-wise model" refer to model 1.

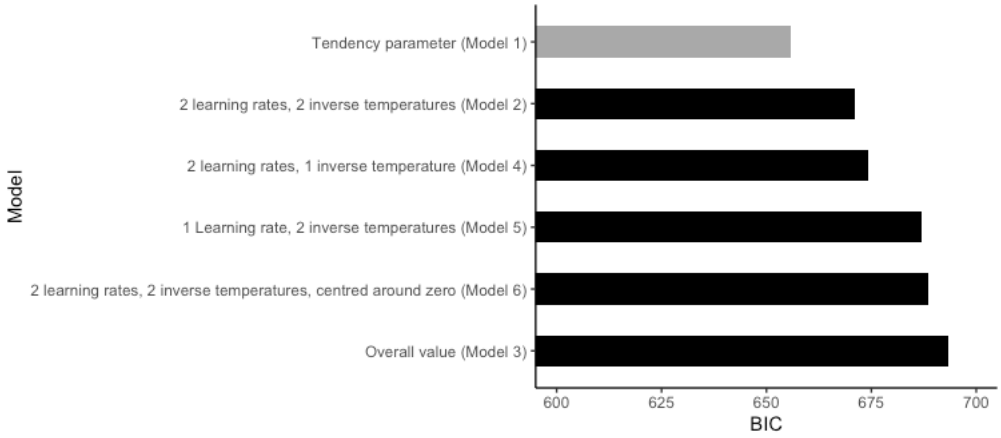

Figure 5.1. Model comparison of the block-wise models for the sample of participants with low mood. As in the previous study, model 1 fitted the data best.

5.2.3.2 Constant inverse temperature models

In addition, we used the constant inverse temperature approach described above to obtain a more robust estimate of the inverse temperature. Model comparison was performed across the same six models as in the previous section, with the only difference being that the inverse temperature parameters in each model were kept constant across the whole session.

Since the inverse temperature parameters were fitted across the whole session, parameter estimation had to be performed for all parameters for the whole session at the same time (omitting the first 10 trials of each block). This resulted in up to 19 parameters (depending on the model) that had to be estimated in parallel. Since calculating the entire posterior distribution for this number of parameters would be too computationally expensive, an alternative method of parameter estimation had to be applied. Parameters were estimated in STAN (Stan Developement Team, 2022) which uses a Markov Chain Monte Carlo (Betancourt & Girolami, 2015) algorithm to sample from the Bayesian posterior distribution. 4 chains were run for 10,000 iterations (adapt_delta = 0.99, max_treedepth = 12). The Gelman-Rubin statistics (R-hat) was below 1.1 for all estimates, indicating good convergence (Vehtari, Gelman, Simpson, Carpenter, & Bürkner, 2021). The prior mean was set to 0 for the inverse temperature (log-transformed), and to -0.5 for learning rates (quantile-transformed) which is the average value observed in the block-wise model. The standard deviation for priors was set to 1. STAN models were run on the University of Oxford Advanced Research Computing (ARC) cluster.

Two models fit the data almost equally well (Figure 5.2). The best fitting model was model 6 with two learning rates and two inverse temperature parameters (one for win and one for loss outcomes)(14 parameters). The second-best model was model 4 with two learning rates and only one inverse temperature (13 parameters).

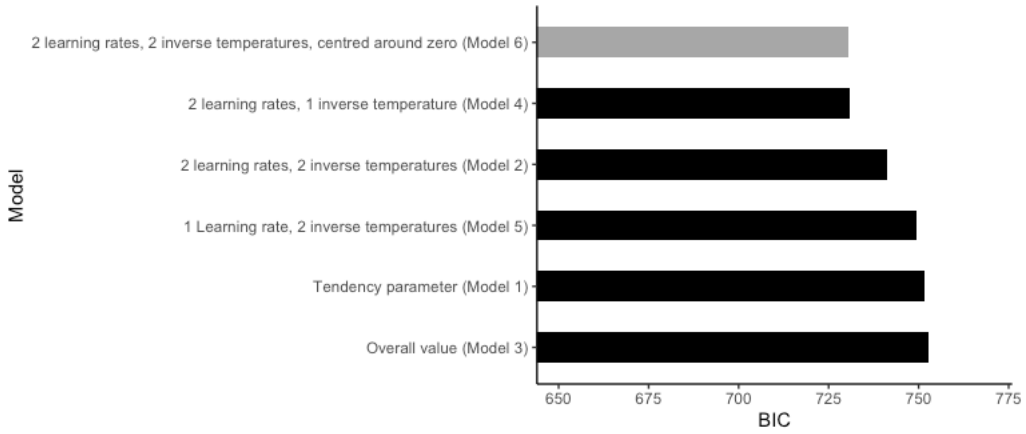

Figure 5.2. Model comparison for the constant inverse temperature models. Model 6 was the best fitting model, although model 4 fitted the data almost equally well.

To understand why model 4 and 6 provided a similarly good fit, model recovery was performed for these two models. The estimated parameter values from model 4 and model 6 for each participant and each session were used to simulate choices in the task. Model 4 and model 6 were then fitted to each of the 330 datasets generated from model 4 and 6. For each dataset, the model that better fits the data was identified using the Bayesian Information Criterion (BIC). Model recovery was evaluated by calculating the percentage of datasets for which the model that generated the data also fitted the data best. The resulting confusion matrix shows that datasets generated from model 4 were best fit by model 4 in 93% (Figure 5.3A). Datasets generated from model 6, however, were only best fit by model 6 in 40%. The reason for that was that the inverse temperature parameters for wins and losses were highly correlated and for roughly half of the datasets were almost identical (Figure 5.3C). In these cases, model 4 with only one inverse temperature parameter resulted in a similar log likelihood but included a smaller number of parameters, thus yielding a lower BIC. To analyse the likelihood of the observed datasets being generated from either of the two models, the inverse confusion matrix was calculated using Bayes' rule (Figure 5.3B). The results suggest that given that model 6 provided the best overall fit, the likelihood that the data were actually generated from model 6 was 85%.

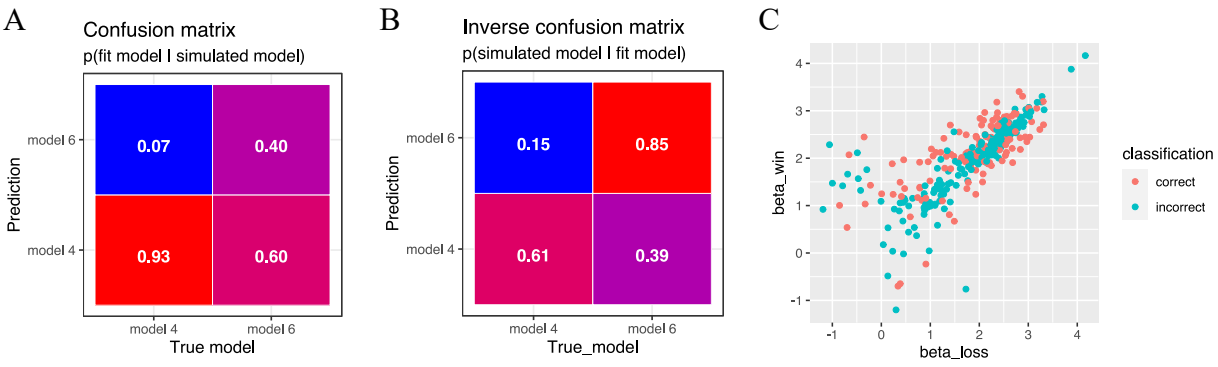

Figure 5.3. Model recovery for model 4 and 6 which fitted the data almost equally well. (A) shows the confusion matrix, i.e. the probability that a dataset simulated from a given model (True model) is best fit by the true or the other model respectively (Prediction). (B) shows the inverse of the confusion matrix, i.e. the probability that a given dataset that is best fit by one of the models (Prediction) is generated from the same or the other model respectively (True model). (C) shows the correct and incorrect classification of datasets generated from model 6 based on the inverse temperature parameters. Datasets with similar inverse temperature for wins and losses were better fit by model 4. Incorrect classifications were also common for low inverse temperatures.

We decided to use the parameter estimates from model 6 for three reasons. First, model 6 provided a (slightly) better model fit. Second, the inverse confusion matrix suggests that given that model 6 fits the data better, the data are very likely generated from model 6. Third, some participants' behaviour was best fit by differential inverse temperatures for wins and losses. There is no negative effect of adding a second inverse temperature parameter. If a dataset is better described by a model using one inverse temperature, the two inverse temperature estimates should be identical (or very similar), without any effect on the learning rate estimates. All subsequent analyses for the constant inverse temperature model refer to model 6.

The simulations run for model recovery were also used for parameter recovery. Estimated and actual parameters were highly correlated for model 6, with correlation coefficients between .73 to .86 (Figure 5.4).

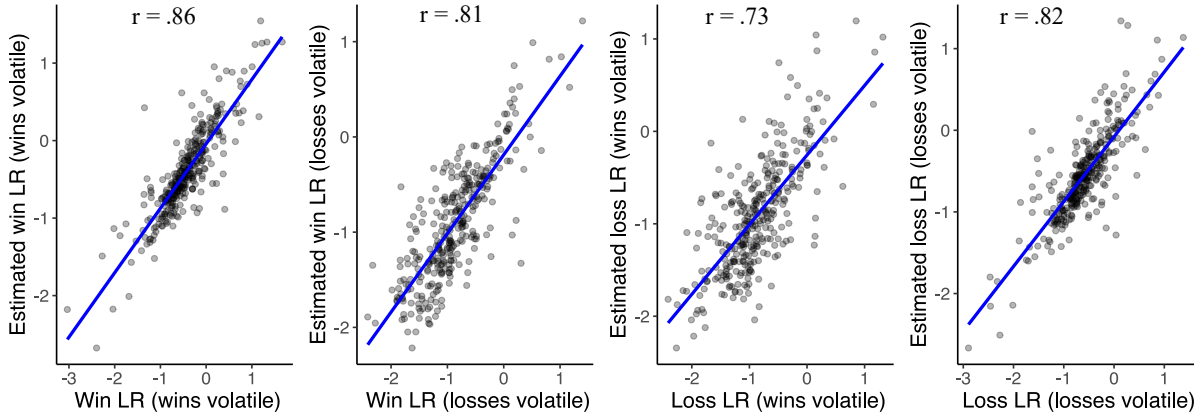

Figure 5.4. Parameter recovery for model 6 for the parameters of interest, i.e. win and loss learning rates in the two blocks during stimulation (block 2 and 3, wins-volatile and losses-volatile respectively). All parameter estimates were highly correlated with the original parameter values. Parameter recovery was better for learning rates of the volatile outcome (win learning rate in the wins-volatile block, and loss learning rate in the losses-volatile block).

5.2.4 Statistical analysis

We started by analysing non-computational measures which included total winnings, number of wins received, number of losses received, win-stay probability (probability of choosing the same option on the next trial after receiving a win outcome), no-loss-stay probability (probability of choosing the same option on the next trial after not receiving a loss outcome), loss-switch probability (probability of choosing the other option on the next trial after receiving a loss outcome) and no-win-switch probability (probability of choosing the other option on the next trial after not receiving a win outcome).

As a non-computational measure of positive bias, we included the proportion of win-driven choices. This measure was based on choices that participants made after a trial in which they received both a win and a loss outcome in response to their chosen shape. These trials are of particular interest because while they result in no net financial change, they can differentiate between two opposite strategies: if participants are primarily trying to maximise the number of wins they receive, then on the next trial they should once again choose the same shape, because it was associated with the win (both outcomes) on the previous trial. However, if participants are mainly trying to avoid losses, on the next trial they should instead pick the alternative shape that was not associated with the loss (neither outcome) on the previous trial. The proportion of "win-driven choices" is therefore the number of such trials in which participants again choose the shape that was associated with both outcomes, divided by the total number of trials in which the win and loss were associated with the same shape.

Regarding computational parameters, learning rates and inverse temperature estimates were analysed. Depression has been associated with difficulties adjusting learning rates to volatility (Browning et al., 2015; Gagne et al., 2020). To quantify this, we calculated the relative adjustment of win and loss learning rates between volatile versus stable blocks. *Win learning rate adjustment* was defined as the difference between the win learning rate in the wins-volatile block and the win

learning rate in the losses-volatile block. *Loss learning rate adjustment* was defined as the difference between the loss learning rate in the losses-volatile block and the loss-learning rate in the wins-volatile block. To capture the extent to which learning rate adjustment was biased towards either win or loss outcomes, a measure for *learning rate adjustment bias* was calculated, which was defined as loss learning rate adjustment minus win learning rate adjustment.

All analyses were performed in RStudio (Version 1.4.1717, R 4.1.1). All dependent variables (apart from stay probabilities) were analysed in mixed ANOVAs, with the main factor of interest being the factor Sample (general population vs. low mood). All analyses on non-computational measures included the between-subject factor Block Order (wins-volatile first vs. losses-volatile first) and the within-subject factors Volatility (both-volatile, wins-volatile and losses-volatile) and Time (first half vs. second half). To test our first hypothesis that participants with low mood would show increased loss learning rates, learning rate estimates were analysed in ANOVAs including all factors mentioned above and the additional key within-subject factor of Valence (win learning rate vs. loss learning rate). Significant two-way interactions of Sample with Valence or Volatility and three-way interactions between Sample, Valence and Volatility were followed up. To test our second hypothesis that low mood would be characterised by decreased learning rate adjustment, learning rate adjustment was analysed in an ANOVA including the factors mentioned above apart from Volatility (since learning rate adjustment is the difference in learning rates between the Volatility conditions). Significant interactions between Sample and Valence were followed up. To test whether low mood was characterised by a relative bias in adjusting win versus loss learning rates, the *learning rate adjustment bias* was analysed in an ANOVA including all factors mentioned above apart from Volatility and Valence. Win-stay probability and no-loss-stay probability were analysed with Wilcoxon rank-sum tests due to violations of the normality assumption.

All computational and non-computational measures were correlated to the BDI and trait anxiety score. Significance of correlations was assessed using t-tests.

All analyses were repeated after removing outliers. A datapoint was identified as potential outlier if it was more than 1.5 times the interquartile range below the first or above the third quartile. For each outcome measure, outliers were removed separately for the levels of the factors of interest (i.e. separately for the general population and low mood sample, and win and loss outcomes where appropriate). Statistics are reported for the entire dataset unless outlier removal had an impact on the results (for completeness, figures including outliers are reported in the supplementary analyses section for all analyses that outlier removal had an impact on). P-values below 0.05 are considered significant, whereas p-values between 0.05 and 0.10 are interpreted as trends.

5.3 Results

5.3.1 Effects of low mood on non-computational measures

As expected, participants in the "low mood" sample showed significantly higher depression and trait anxiety scores than participants in the general population sample (BDI: $t(73.8) = 12$, $p < .001$; trait anxiety scale (STAI): $t(79) = 11.7$, $p < .001$ (two-sample Welch t-test))(Figure 5.5A and 5.5B). BDI and trait anxiety scores were highly correlated within and across groups (across groups: $r = .90$, $t(80) = 19.3$, $p < .001$, within groups: low mood: $r = .81$, $t(41) = 8.9$, $p < .001$, general population: $r = .62$, $t(37) = 4.8$, $p < .001$)(Figure 5.5C).

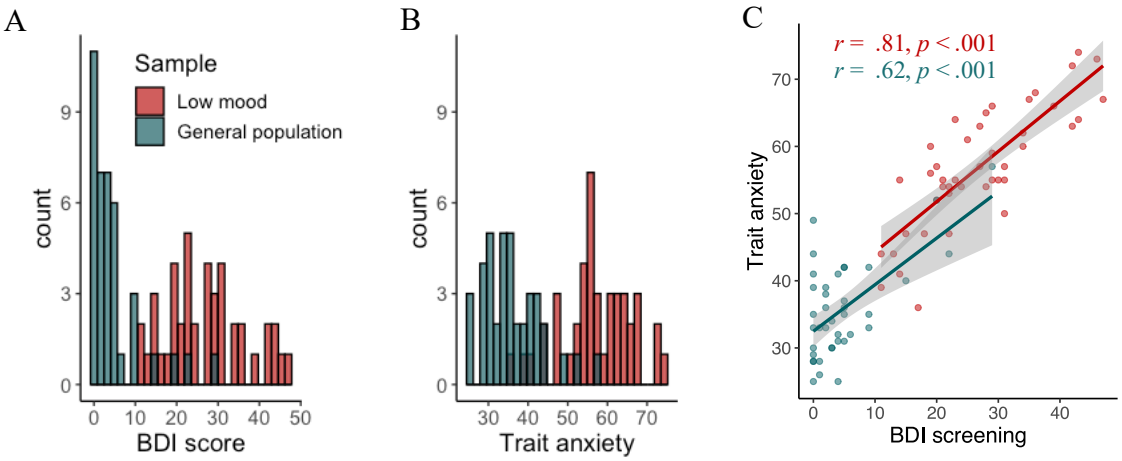

Figure 5.5. Distribution of (A) BDI scores and (B) trait anxiety scores (State-trait Anxiety Inventory (STAI), (Spielberger, 1983)) for the "low mood" and "general population" samples. (C) BDI and trait anxiety scores were highly correlated. BDI scores can range from 0 to 63 with higher scores indicating more severe depressive symptoms. Trait anxiety scores can range from 20 to 80.

Overall, participants suffering from low mood won less money in the task ($F(1,79) = 8.3$, $p = .004$), received fewer wins ($F(1,79) = 8.0$, $p = 005$), and more losses (Sample: $F(1,79) = 4.5$, $p = .036$)(Figure 5.6).

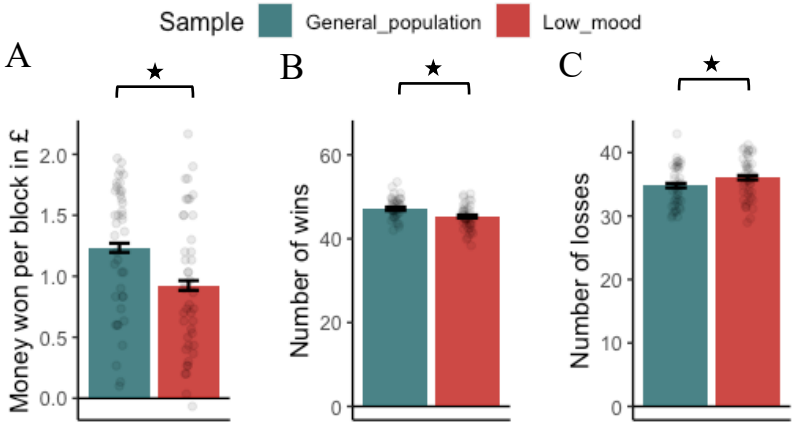

Figure 5.6. Effect of low mood on task performance. Participants in the low mood sample won less money (A) and received less wins (B) and more losses (C) compared to the general population sample.

Participants in the low mood sample switched more often ($F(1,79) = 5.0$, $p = .02$), were less likely to stay after receiving a win (win-stay probability: $W = 1118$, $p = .018$) or after not receiving a loss (no-loss-stay probability: $W = 1136$, $p = .012$). There was no effect of Sample on the no-win-switch probability ($F(1,79) = 2.7$, $p = .10$) but a trend towards higher loss-switch probability in the low mood sample ($F(1,78) = 3.2$, $p = .076$)(Figure 5.7).

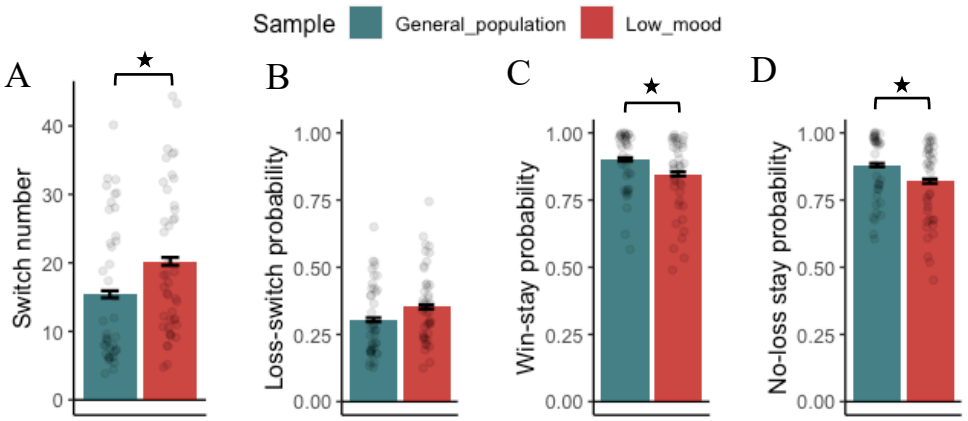

Figure 5.7. Effect of low mood on switch behaviour. Participants in the low mood sample switched more often (A), showed a trend towards higher loss-switch probability (B), had lower win-stay probabilities (C), and lower no-loss-stay probabilities (D).

Participants with low mood did not differ from the general population sample in the proportion of win-driven choices (Sample: $F(1,79) = 0.02$, $p = .87$).

Across groups, BDI scores were negatively correlated with total winnings ($r = -.23$, $t(79) = 2.0$, $p = .042$) and the number of received wins ($r = -.35$, $t(77) = -3.3$, $p = .001$). However, within the two samples no significant relationship between BDI or trait anxiety scores with any of the non-computational measures was observed (all $p > .26$).

5.3.2 Effects of low mood on computational measures

5.3.2.1 Effect of low mood in the block-wise model

To test our first hypothesis that individuals with low mood would show increased punishment vs. reward learning, learning rate parameters derived from model 1 were compared between the two samples. The low mood sample showed comparable learning rates to the general population sample (Sample: $F(1,79) = 1.9$, $p = .16$; Sample x Valence: $F(1,79) = 0.1$, $p = .66$); Sample x Valence x Volatility: $F(2,158) = 0.1$, $p = .86$)(Figure 5.8A).

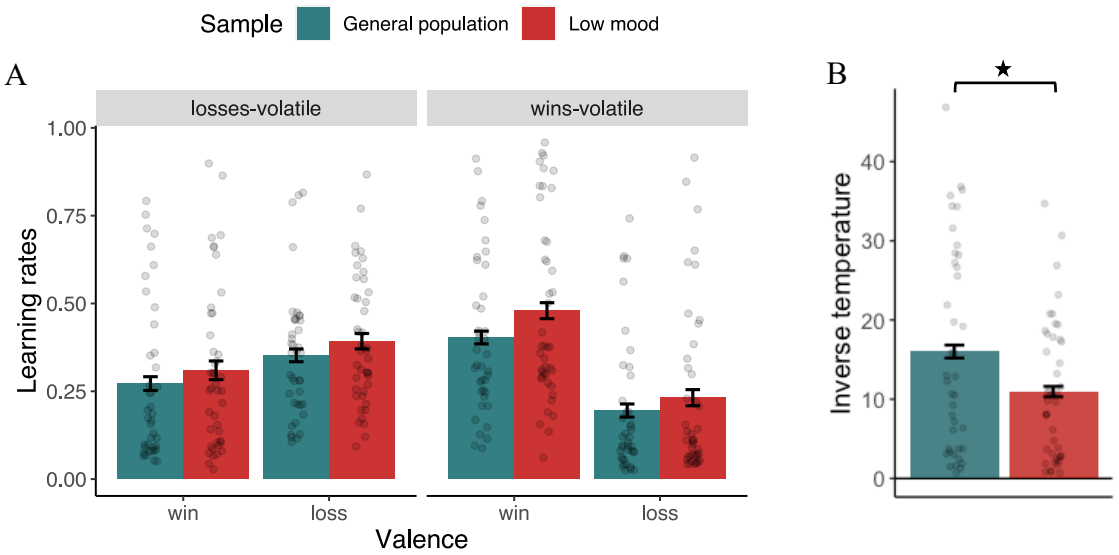

Figure 5.8. Effect of low mood on learning rates and inverse temperature in the block-wise model. The low mood sample had comparable win and loss learning rates (A), but lower inverse temperature estimates (B) compared to the general population sample.

To test our second hypothesis that individuals with low mood adjust their learning rates in response to volatility to a smaller extent than healthy individuals, an ANOVA on learning rate adjustment was conducted. There was no main effect of Sample ($F(1,79) = 0.1$, $p = .65$) or interaction effect between Sample and Valence ($F(1,79) = 0.9$, $p = .34$). The learning rate adjustment bias measure

(i.e. loss learning rate adjustment minus win learning rate adjustment) did not differ either between the two samples ($F(1,79) = 0.9$, $p = .34$)(Figure 5.9).

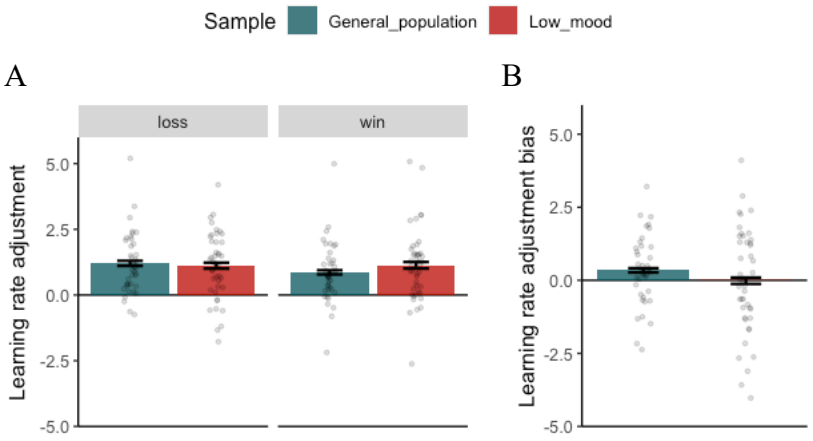

Figure 5.9. Effect of low mood on learning rate adjustment. Participants suffering from low mood did not differ from the general population sample in learning rate adjustment (A) or learning rate adjustment bias (B).

To test whether individuals with low mood differed from healthy individuals in choice randomness, the inverse temperature estimates were compared between the two samples. Individuals with low mood had lower inverse temperature estimates, i.e. were more random in their choice behaviour (Sample: F(1,79) = 4.9, p = .029)(Figure 5.8B).

No computational measure was significantly correlated with the BDI or trait anxiety score (all $p > .11$).

5.3.2.2 Effects of low mood in the constant inverse temperature model

Next, we ran the same analyses on the estimates derived from the constant inverse temperature model. To test our first hypothesis that individuals with low mood would show increased punishment vs. reward learning, learning rate estimates were analysed in an ANOVA with Sample and Valence as the factors of interest. Contrary to this hypothesis, there was no main effect of

Sample (Sample: $F(1,77) = 0.5$, $p = .46$). There was a trend towards an interaction between Sample and Valence (Sample x Valence: $F(1,77) = 3.2$, $p = .074$), but the effect of Sample was neither significant for win nor for loss learning rates (win learning rate: $F(1,77) = 0.03$, $p = .86$); loss learning rate: $F(1,77) = 2.4$, $p = .12$) (Figure 5.10A). Win and loss learning rates were not significantly correlated with the BDI or trait anxiety score (all $p > .29$)

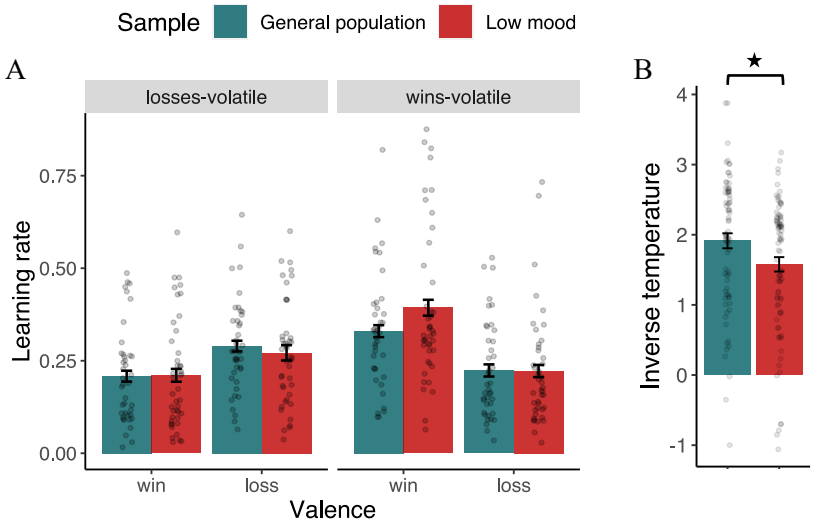

Figure 5.10. Effect of low mood on learning rates and inverse temperature in the constant inverse temperature model. (A) No difference in learning rates between participants with low mood and the general population. (B) Participants with low mood showed lower inverse temperatures independent of valence (for wins and losses).

To test our second hypothesis that individuals with low mood adjust their learning rates to a smaller extent than healthy individuals, learning rate adjustment was analysed in an ANOVA with Sample and Valence as main factors of interest. There was a Sample x Valence interaction (after outlier removal: $F(1,72) = 4.4$, $p = .038$ (see Supplementary Figure 5.13); before outlier removal: $F(1,79) = 4.0$, $p = .046$). Post-hoc tests indicated that there was no significant main effect of Sample on win learning rate adjustment ($F(1,72) = 1.3$, $p = .24$) but a trend towards lower loss learning rate adjustment in the sample with low mood ($F(1,72) = 3.2$, $p = .076$)(Figure 5.11A). However, in line with the interaction effect of Sample and Valence on learning rate adjustment, there was a main effect of Sample on the combined measure of *learning rate adjustment bias* (after outlier removal: $F(1,72) = 4.4$, $p = .038$; before outlier removal: $F(1,79) = 4.0$, $p = .046$)(Figure

5.11B). Individuals with low mood showed a lower *learning rate adjustment bias*. That is, while the general population adjusted their win and loss learning rates to a similar extent (one-sample t-test against zero of the learning rate adjustment bias (general population): $t(35) = -1.3$, $p = 0.17$), individuals with low mood adjusted their loss learning rate significantly less than their win learning rate (one-sample t-test against zero of the learning rate adjustment bias: $t(39) = -3.2$, $p = .002$).

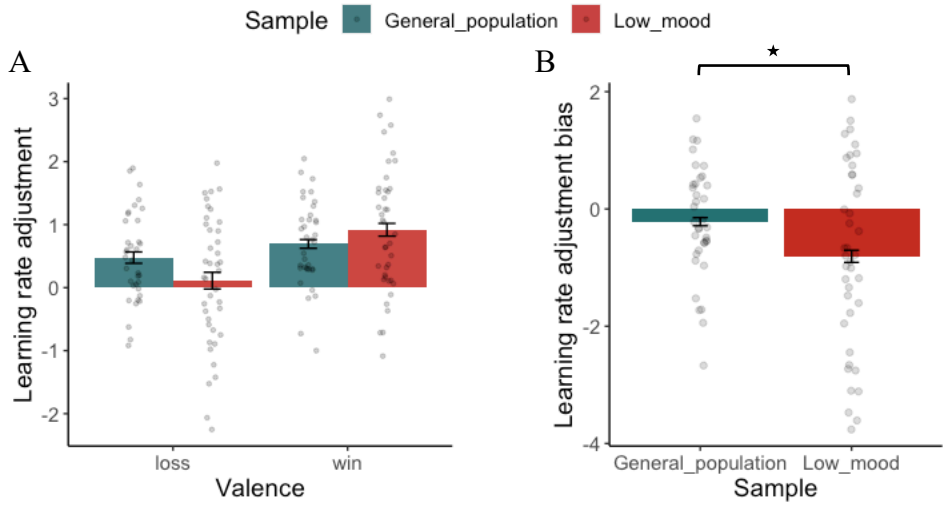

Figure 5.11. Effect of Sample on learning rate adjustment. (A) Interaction effect between Sample and Valence on learning rate adjustment. Participants with low mood showed a trend towards lower loss and higher win learning rate adjustment. (B) In line with this, participants with low mood showed a lower learning rate adjustment bias.

To test for a linear relationship between mood and learning rate adjustment, BDI and trait anxiety scores were examined for correlation with learning rate adjustment. Across groups, BDI scores correlated negatively with the learning rate adjustment bias, i.e. individuals with more severe depressive symptoms showed a more negative *learning rate adjustment bias* ($r = -.25$, $t(75) = -2.2$, $p = .027$). This is in line with the group differences observed in the ANOVA. In line with this, there was a non-significant negative correlation between BDI scores and loss learning rate adjustment ($r = -.18$, $t(75) = -1.5$, $p = .11$) and a trend towards a positive correlation between BDI scores and win learning rate adjustment ($r = .20$, $t(73) = 1.8$, $p = .070$). However, the same correlations conducted within each group were not significant (see Figure 5.12), indicating that the

correlation between BDI and learning rate adjustment was driven by the group difference. On a descriptive level, the correlations in the low mood sample (which has a higher variance in the BDI scores than the general population sample) were in the same direction as the correlations across the two samples.

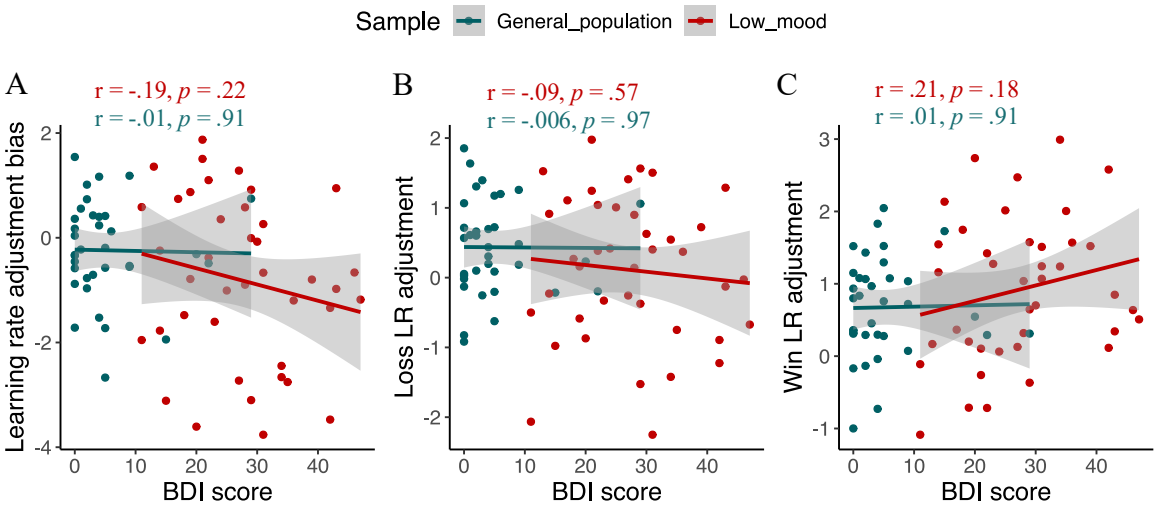

Figure 5.12. Correlations between BDI score and learning rate adjustment bias (A), loss learning rate adjustment (B) and win learning rate adjustment (C).

To test whether individuals with low mood differed from healthy individuals in choice randomness, the inverse temperature estimates for wins and losses were analysed in an ANOVA with Sample and Valence as factors of interest. There was a marginally significant main effect of Sample on the inverse temperature estimates ($F(1,78) = 3.9$, $p = .050$). Participants with low mood showed lower inverse temperatures for wins and losses, independent of their valence (Sample x Valence: $F(1,78) = 0.17$, $p = .67$)(Figure 5.10B). We tested for correlations of the BDI or trait anxiety score with the inverse temperature estimates for wins, losses or averaged across both outcomes. Across groups, there was a significant negative correlation between the BDI score and win inverse temperature ($r = -.21$, $t(80) = -1.9$, $p = .047$) which became non-significant after removing the most extreme datapoint (see Supplementary Figure 5.14). There was no significant correlation between BDI or trait anxiety score with the loss inverse temperature or average inverse temperature, neither across nor within groups (all $p > .13$).

To summarise the overall findings, the non-computational measures indicated that individuals with low mood earned less money in the task, were more likely to switch to the other option and were less likely to stay with the previous option. In the block-wise model, the only difference observed between the two samples was a decreased inverse temperature for participants with low mood. In the constant inverse temperature model, low mood was associated with decreased *learning rate adjustment bias*. While healthy individuals adjusted win and loss learning rates to the same extent, individuals with low mood adjusted their loss learning rate less than their win learning rate. Decreased inverse temperature parameters in individuals with low mood were also observed in the constant inverse temperature model.

5.4 Discussion

The goal of this chapter was to investigate negative biases in reinforcement learning in depression. Participants with low mood and healthy participants performed a task in which the relative informativeness of positive and negative outcomes was manipulated. In contrast to our first hypothesis, participants with low mood did not show increased learning rates from negative outcomes. However, in line with our second hypothesis low mood was associated with altered adjustment of learning rates. While healthy individuals adjusted win and loss learning rates to the same extent to changes in informativeness, participants with low mood adjusted their loss learning rates less than their win learning rates. However, this finding was only observed in one of the two modelling approaches.

A current hypothesis states that negative biases might arise from increased punishment learning rates (Aylward et al., 2019; Pike & Robinson, 2022). We therefore hypothesised that participants with low mood would show higher loss learning rates than healthy individuals. In contrast, we did not find evidence for increased loss learning rates in either of the two models. Evidence regarding the hypothesis of increased punishment learning in depression is mixed since some studies failed to detect valence-specific effects of depressive symptoms on learning rates (Beevers et al., 2013;

Kunisato et al., 2012; Mukherjee et al., 2020). Most direct evidence supporting the hypothesis of increased punishment learning rates comes from a study using a 4-armed bandit task (Aylward et al., 2019). Similar to our paradigm, the four options were independently associated with positive and negative outcomes. In contrast to our study, the reward and punishment probabilities drifted continuously over time. The study found that participants with depressive symptoms showed a selective increase in punishment learning rates compared to healthy participants. The absence of this effect in our study might be caused by differences in the experimental paradigm. Our manipulation of the relative informativeness of rewards and punishments had a very strong effect on learning rates which might have predominated over a potential smaller effect of mood.

In line with our second hypothesis, we found evidence that participants with low mood adjusted their learning rates in a different way to changes in informativeness than healthy participants. While healthy participants adjusted their win and loss learning rates to an equal extent, participants with low mood adjusted their loss learning rate less than their win learning rate. Furthermore, and in line with this, the *learning rate adjustment bias* was negatively correlated with the severity of depressive symptoms across groups (non-significant negative correlation within the low mood group), i.e. individuals with more severe depressive symptoms showed a relatively smaller adjustment of loss learning rates. The difference in *learning rate adjustment bias* seems to be caused both by decreased loss learning rate adjustment as well as by increased win learning rate adjustment. Neither win or loss learning rate adjustment on their own were significantly associated with low mood. While the group difference seemed to be mainly driven by differences in loss learning rate adjustment, win learning rate adjustment showed a trend towards a positive correlation with severity of depressive symptoms across groups (the positive correlation within the low mood group was not significant but showed a similar strength).

Depression and anxiety have previously been associated with deficits in learning rate adjustment (Browning et al., 2015; Gagne et al., 2020). Adjustment of reward learning rates has been found to be decreased (Gagne et al., 2020) or unaffected (Browning et al., 2015). The increase in win

learning rates observed in this study therefore contradicts previous studies. In contrast to previous studies, in our paradigm, participants were simultaneously presented with positive and negative outcomes, the informativeness of which was manipulated independently. Tracking two types of outcomes simultaneously requires a lot of cognitive resources, such as attention and working memory. Participants therefore need to decide how to allocate their cognitive resources to reward and punishment learning. A potential explanation of our findings might be that individuals with low mood allocated more cognitive resources to positive than to negative outcomes while healthy individuals allocated the same amount of cognitive resources to positive and negative outcomes. However, this finding seems counter-intuitive, since depression is associated with increased processing of punishments relative to rewards (Beevers et al., 2013; Eshel & Roiser, 2010). One potential explanation might be that participants have difficulties tracking the informativeness of losses and therefore spend more resources on tracking win outcomes as a compensatory strategy. It is unclear how the changes in learning rate adjustment in participants with low mood affected task performance, i.e. whether increased win learning rate adjustment was adaptive. Participants with low mood generally performed worse which is most likely caused by their significantly increased choice randomness (lower inverse temperature). The effect of altered learning rate adjustment on task performance could be de-coupled from increased choice randomness by simulating choices for both samples with the observed learning rate estimates, but equivalent inverse temperature.

In line with previous research (Kunisato et al., 2012; Mukherjee et al., 2020; Pike & Robinson, 2022), the sample of participants with low mood had lower inverse temperature estimates, i.e. their choices were less influenced by the estimated outcome probabilities. This might be caused by lower levels of concentration and motivation which are commonly affected in depression (Grahek, Shenhav, Musslick, Krebs, & Koster, 2019; Rock, Roiser, Riedel, & Blackwell, 2014). Alternatively, decreased inverse temperature estimates can indicate that participants employed a strategy which was not captured by the model. In this case the model fit should be worse for participants with low mood. To test this, the BIC was compared between samples. The model fit for both types of models was indeed worse for the sample with low mood as indicated by a higher

BIC (however, the BIC was not significantly correlated with BDI scores; see Supplementary Figure 5.16 and Supplementary Figure 5.17). Therefore, it is possible that individuals with low mood applied a different strategy in the task compared to healthy individuals.

We analysed this dataset with two different computational modelling approaches which yielded different results. In addition to the block-wise modelling approach applied in our previous study, we used a modelling approach with constant inverse temperature parameters, since the volatility manipulation did not affect the inverse temperature in previous studies. This approach is based on the assumption that choice randomness is constant across all task blocks. To test this assumption, we compared inverse temperature estimates from the block-wise model between the different task blocks. The inverse temperature actually differed between the task blocks, with lower values in the losses-volatile condition (Supplementary Figure 5.15). This indicates that the constant inverse temperature model might be over-simplified. While the effect of low mood on the inverse temperature was present in both models, effects on the learning rates were only observed in the constant inverse temperature model, which calls the robustness of the observed effects into question. Further research is needed to test whether this finding replicates in an independent dataset. Another limitation is that we did not correct for multiple comparisons which increases the risk for false-positive findings. However, similar results have been observed in previous studies (Browning et al., 2015; Gagne et al., 2020; Pike & Robinson, 2022) which increases our confidence that these effects would replicate.

Our findings may have clinical implications. Since increased processing of negative information is hypothesised to causally contribute to depressive symptoms, cognitive interventions often aim at shifting information processing away from negative towards positive information. For example, attentional bias modification aims at biasing attention towards positive stimuli by consistently pairing a target with positive stimuli (Dai, Hu, & Feng, 2019; MacLeod, Rutherford, Campbell, Ebsworthy, & Holker, 2002). Our results suggest that rather than training participants to generally attend to positive rather than negative information, interventions should perhaps instead train

participants in adjusting their behaviour to changes in informativeness. More precisely, we suggest that a training task should teach participants to keep track of the informativeness of positive and negative outcomes and preferentially learn from the outcome which is currently more informative. While our results are purely correlational and do not allow for causal conclusions, future research could test whether training in learning rate adjustment might generalise to other domains and might ultimately improve mood. With respect to the effect of tDCS observed in our previous study in healthy volunteers (increased reward learning), the current findings suggest that inducing a general increase in reward learning might not actually counter-act altered reinforcement learning in depression. Rather, according to our findings in this chapter, the effect of tDCS on learning rate *adjustment* would be of particular interest.

5.5 Conclusions

To summarise, the aim of this project was to assess alterations in reward and punishment learning associated with depression. Healthy individuals and individuals with low mood performed a task in which the relative informativeness of positive and negative outcomes was manipulated. In contrast to our hypothesis and previous work (Aylward et al., 2019; Pike & Robinson, 2022), we did not observe evidence for increased punishment learning rates in low mood. However, we found that low mood was associated with alterations in learning rate adjustment. While healthy individuals adjusted their win and loss learning rates to the same extent, individuals with low mood adjusted their loss learning rate less than their win learning rate. In our previous study (chapter 4), we found that bifrontal tDCS increased reward learning rates in healthy individuals. Based on the hypothesis that negative biases might be caused by increased punishment learning rates we had reasoned that this effect might be useful to counter-act negative biases in depression. However, based on the findings in this chapter, it is unclear whether an increase in reward learning would be beneficial. Our next study (chapter 6) investigated the effect of bifrontal tDCS in participants with depressive symptoms. The current findings indicate that the effect of tDCS on learning rate *adjustment* in that study might be of particular interest.

5.6 Supplementary analyses

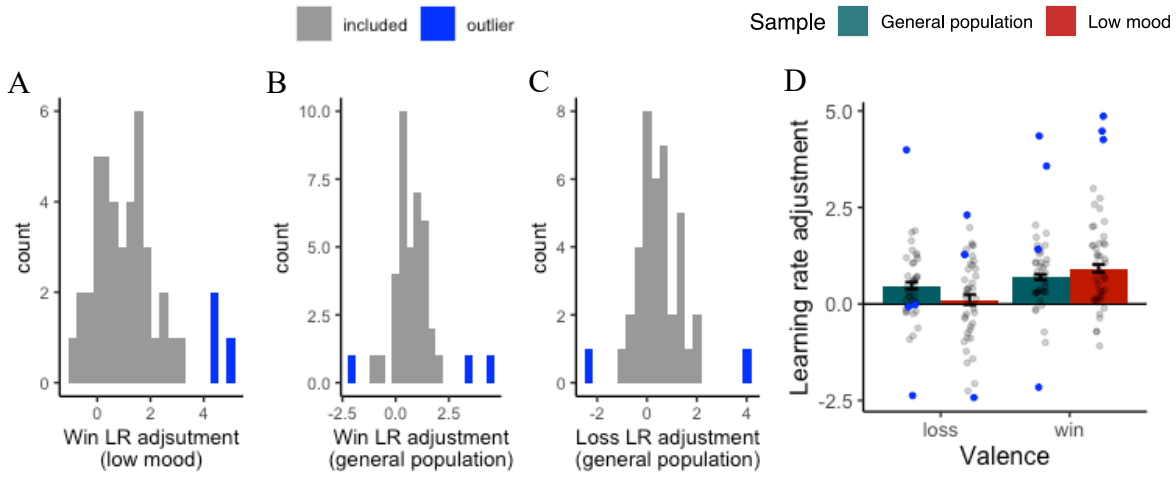

Figure 5.13. Outlier removal for learning rate adjustment. Outliers were identified separately for win and loss learning rate adjustment, and for the low mood and general populations sample (low mood: 3 outliers, general population: 4 outliers). (A) Outliers for win learning rate adjustment in the low mood sample, (B) outliers for win learning rate adjustment in the general population sample, (C) outliers for loss learning rate adjustment in the general population sample (there were no outliers for loss learning rate adjustment in the low mood sample). (D) shows the difference in learning rate adjustment between the samples reported in the main text. Outliers are overlaid in blue.

Before outlier removal: r = -0.21, $t(80) = -2.0$, $p = .047$
After outlier removal: r = -0.14, $t(79) = -1.3$, $p = .189$

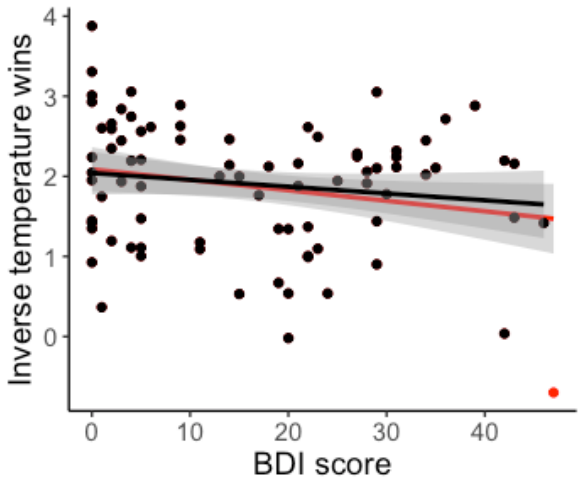

Figure 5.14. Correlation of BDI score with inverse temperature for wins before and after outlier removal (red dot = outlier).

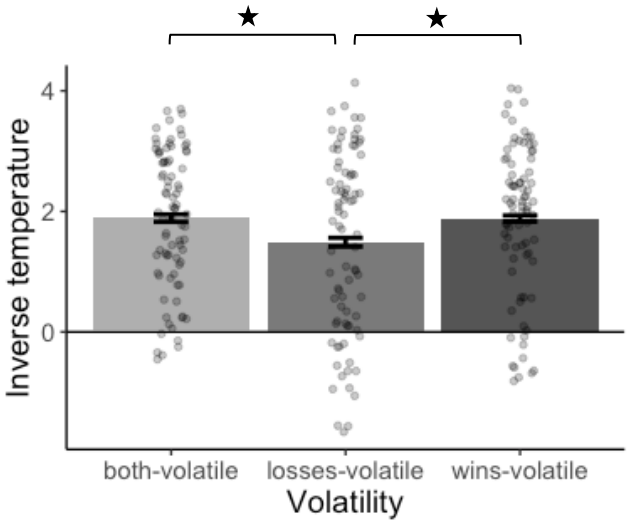

Figure 5.15 Differences in inverse temperature between volatility conditions in the block-wise model. An ANOVA with the inverse temperature estimates as dependent variable and Volatility, Time and Sample as independent variables was conducted. There was a significant main effect of Volatility ($F(2,162) = 9.0$, $p < .001$). Inverse temperature estimates for the losses-volatile condition were significantly lower than for the both-volatile ($t(82) = 3.3$, $p = .001$) and wins-volatile condition ($t(82)=-3.7$, $p < .001$). There was no interaction between Volatility and Sample ($F(2,162) = 0.9$, $p = .40$) indicating that the effect of Volatility on inverse temperature did not differ between the samples.

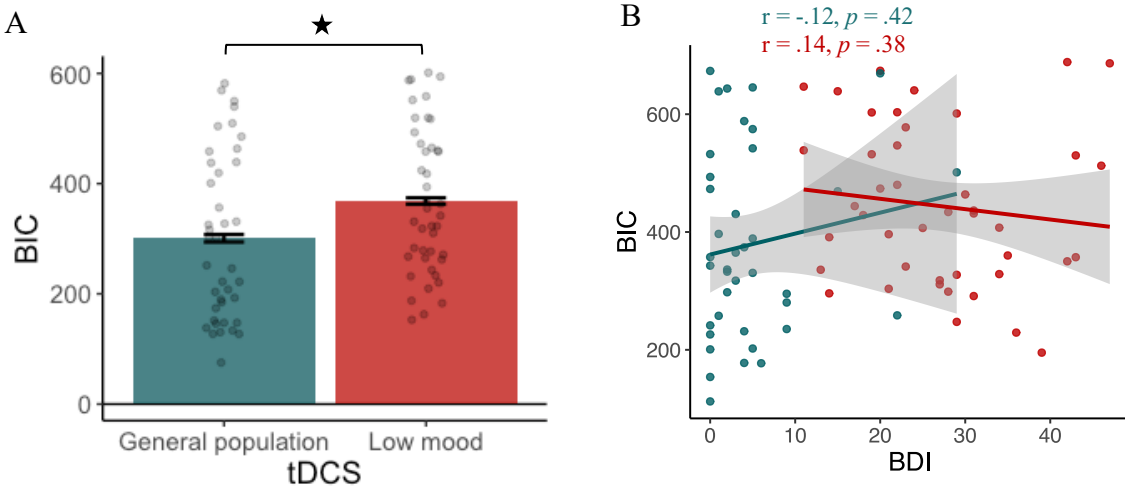

Figure 5.16. Comparison of model fit for the block-wise model between samples. (A) The model fit was worse for participants with low mood as indicated by a higher BIC ($t(77.9) = -2.1$, $p = .033$). (B) The BIC was not significantly correlated with BDI scores.

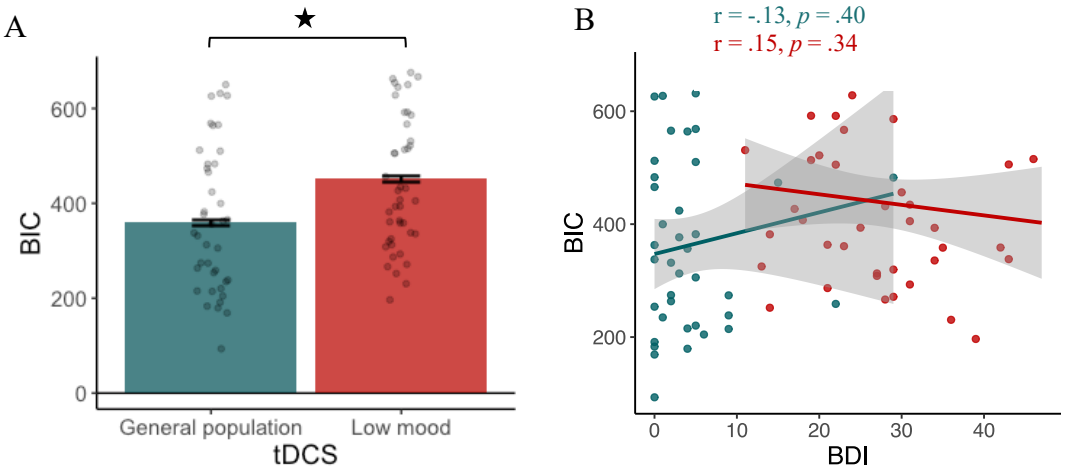

Figure 5.17. Comparison of model fit for the constant inverse temperature model between samples. (A) The model fit was worse for participants with low mood as indicated by a higher BIC ($t(78.7) = -2.3$, $p = .02$). (B) The BIC was not significantly correlated with BDI scores.

6 The effects of bifrontal tDCS on reward learning in low mood

In our first study in healthy volunteers (Chapter 4), we found that bifrontal tDCS, when applied during but not before the performance of a reinforcement learning task, selectively increased reward learning rates. Since negative biases in depression are hypothesised to arise from increased punishment learning, such an increase in reward learning might counter-act negative biases. The aim of the present study was therefore to test whether bifrontal tDCS applied during reinforcement learning could reduce negative biases in individuals with low mood. We hypothesised that online bifrontal tDCS would selectively increase reward learning. In addition, since we found in the previous chapter that low mood was associated with altered learning rate adjustment, we also tested whether tDCS influenced learning rate adjustment. In contrast to our hypothesis, online bifrontal tDCS did not increase reward learning in low mood. However, our findings suggest that online bifrontal tDCS might normalise learning rate adjustment in low mood, by facilitating adjustment of loss compared to win learning rates. This effect was only present when tDCS was applied during as opposed to before task performance, indicating cognitive-state dependency. Since negative biases are hypothesised to result from deficits in adjusting learning rates to informativeness, the effect observed in this study might have clinical benefits in the treatment of depression.

6.1 Introduction

In chapter 4 we found that bifrontal tDCS applied during (but not before) reinforcement learning selectively increased learning from positive outcomes in healthy individuals. Since negative biases in depression are hypothesised to result from increased punishment learning (Aylward et al., 2019; Eshel & Roiser, 2010; Pike & Robinson, 2022), an increase in reward learning might counter-act such negative biases and might thus be beneficial in depression treatment. However, it is unclear whether the effect of bifrontal tDCS we observed in healthy volunteers would generalise to clinical populations. The aim of this study was therefore to investigate the effect of online bifrontal tDCS on reinforcement learning in individuals with depressive symptoms.

If negative biases in depression are caused by increased punishment learning (Aylward et al., 2019; Eshel & Roiser, 2010; Pike & Robinson, 2022), then the increase in reward learning with online bifrontal tDCS we observed in healthy individuals might be beneficial in depression treatment. Since there is also contradicting evidence about whether depression is associated with increased punishment learning (Beevers et al., 2013; Kunisato et al., 2012; Mukherjee et al., 2020), in the previous chapter we tested this in our experimental paradigm. In contrast to our hypothesis, individuals with low mood did not show increased punishment learning. Instead, low mood was associated with altered adjustment of learning rates to the informativeness of positive and negative outcomes. While healthy individuals adjusted win and loss learning rates to the same extent to changes in informativeness, individuals with low mood adjusted their loss learning rate less than their win learning rate. In the treatment of depression, normalisation of this altered learning rate adjustment might therefore be of greater potential benefit than an induced increase in reward learning.

Based on our findings in healthy individuals, we hypothesised that (1) bifrontal tDCS applied during the Information Bias Learning Task would selectively increase reward learning in individuals with depressive symptoms. However, since we found that low mood was associated

with altered learning rate adjustment, we also tested whether bifrontal tDCS might influence learning rate adjustment. This analysis was exploratory as there was no prior evidence suggesting that tDCS might modulate learning rate adjustment. We further hypothesised that (2) the cognitive state during stimulation matters, i.e. bifrontal tDCS applied *before* task performance should not have the same effect as tDCS applied *during* task performance.

6.2 Methods

6.2.1 Sample

85 participants[3] (54 women, mean age = 24.2 years (SD = 4.6 years), see Table 6.1) suffering from low mood completed the study. The advertisement was specifically targeted at individuals currently experiencing low mood. Participants were asked to fill out an online screening including the Beck Depression Inventory II (BDI)(Beck et al., 1996) and a tDCS safety questionnaire. Participants were excluded from the study if they had any contraindication to tDCS, such as medication (apart from the contraceptive pill), neurological conditions, or a family history of epilepsy, metal implants inside the brain, or current pregnancy. Participants who scored at least 10 on the BDI were invited to an online interview during which their history of mental health was assessed in the Structured Clinical Interview for the Diagnostic and Statistical Manual of Mental Disorders (DSM-5)(First, 2014) conducted by a trained PhD student. To ensure participants were not suffering from bipolar disorder, participants were excluded from the study if they met the diagnostic criteria for a past manic or hypomanic episode. The study was approved by the University of Oxford Central University Ethics Committee (R67041/RE002) and all participants gave informed written consent before taking part in the study.

[3] The author of this thesis recruited 75 participants.

Table 6.1. Mean (SD) baseline characteristics for the "online tDCS" and "offline tDCS" groups.

	Online tDCS (n = 41)	Offline tDCS (n = 44)
Sociodemographic data		
Female(%)	24 (58%)	30 (68%)
Age in years (SD)	24.3 (4.8)	24.2 (4.6)
Clinical measures		
STAI-Trait	55.4 (9.0)	57.4 (8.3)
BDI	24.9 (9.1)	27.7 (8.4)

BDI: Beck Depression Inventory-II, score range = 0-63; STAI-Trait: State-Trait Anxiety Inventory (trait form), score range = 20-80.

6.2.2 Experimental design

Participants were assigned to one out of two tDCS conditions. The first 41 participants received tDCS during task performance ("online tDCS" group). To test whether the cognitive state during stimulation is critical, another 44 participants received tDCS at rest ("offline tDCS" group) and performed the task after the stimulation period. All participants attended two experimental sessions in which they received real or sham tDCS in counter-balanced order.

6.2.3 Questionnaires

To assess the level of symptoms of depression and anxiety, the Beck Depression Inventory (BDI-II)(Beck et al., 1996) and the trait-anxiety scale from the State-Trait Anxiety Inventory (STAI)(Spielberger, 1983) were completed at the beginning of the first session. Participants showed a wide range of depression scores ranging from the cut-off score of 10 (minimal depression) to severe depressive symptomatology (Figure 6.1A). Trait anxiety scores also showed a wide range and were highly correlated to the BDI scores ($r = .67$, $t(83) = 8.2$, $p < .001$)(Figure 6.1B and 6.1C). To test for potential changes in affect over the course of the testing sessions, the Positive and Negative Affect Schedule (PANAS)(Watson et al., 1988) and the state anxiety scale from the STAI were completed before and after task performance. Analyses of the questionnaire data are included in the supplementary analysis section.

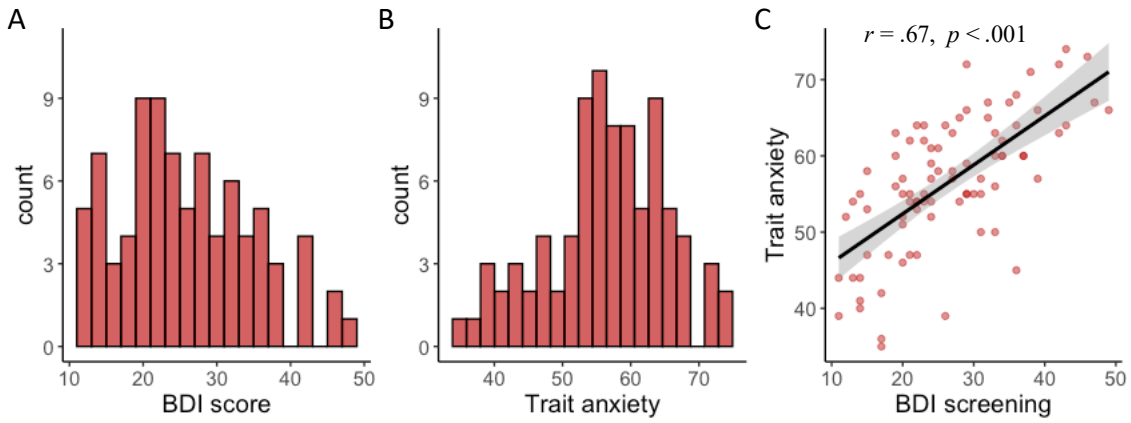

Figure 6.1. Distribution of scores on the BDI (A) and the STAI trait scale (B). BDI and trait anxiety scores were highly correlated (C).

6.2.4 tDCS protocol

The stimulation was delivered through a battery-powered DC-Stimulator Plus (neuroConn). All participants took part in two testing sessions in which they received real or sham tDCS in counter-balanced order. For both types of tDCS, the protocol included a ramp-up and ramp-down period of 10s. Real tDCS was applied for 20 minutes at an intensity of 2mA. For sham stimulation, tDCS was applied at 2mA for 30 seconds, after which the current was ramped down. During the remaining stimulation period, short pulses (110 µA over 15ms) were delivered every 550ms which does not have any physiological effects but enables impedance control. The type of stimulation delivered was determined by a code that was typed into the device to enable double-blinding.

At the start of the tDCS setup, head measurements were taken to locate the vertex. An EEG cap was used to mark the F3 and F4 electrode position (10-20 system) which are an approximation for the left and right DPLFC. The anode was placed on the left, and the cathode on the right DLPFC. The marked areas of the scalp were cleaned using alcohol and high-chloride abrasive electrolyte gel (Abralyt HiCl, Easycap). A 5mm thick layer of gel was applied to the rubber electrodes which were then fixed to the scalp using elastic bands. The impedance was kept below 10µΩ.

6.2.5 Experimental procedure

The experimental procedure was largely the same as in our previous study described in chapter 4. Participants who passed the screening procedure, were invited to attend two testing sessions at the Oxford Centre for Human Brain Activity (OHBA). They were screened for Covid-19 symptoms before entering the building, were asked to sanitise their hands and were given a clinical mask to wear throughout the entire session. After providing informed consent and completing the questionnaires, the tDCS equipment was set up. Participants were then given the instructions for the Information Bias Learning Task (IBLT)(Pulcu & Browning, 2017) which consisted of six blocks of 80 trials (see section 4.2.4). In block 1 and 6, both wins and losses were volatile. Block 2 to 4 were wins-volatile and losses-volatile blocks in alternating order. Participant were asked to perform 10 practice trials to make sure they understood the instructions. After performing the first task block (both-volatile), the stimulation was turned on for 20 minutes. Participants in the "online" condition performed the second and third task block during the stimulation period, and the three remaining blocks after the stimulation had finished. Participants in the "offline" condition were stimulated at rest, and completed the remaining five task blocks afterwards. After task completion, the tDCS setup was removed, and participants were asked to complete another set of questionnaires including the state anxiety scale of the STAI, the PANAS, an adverse effects questionnaire and questions on strategies used during the task.

6.2.6 Computational modelling

The same two different modelling approaches as in the previous chapter were applied. In the first approach, all model parameters were fitted separately to each of the six task blocks ("block-wise" models). In the second approach, the inverse temperature was kept constant across all six task blocks while all other parameters were fitted separately to each block ("constant inverse temperature" models). The first 10 trials per block were omitted in both modelling approaches. For each modelling approach, six different models were fitted to the data and parameter estimates

from the model that provided the best fit were analysed. Model comparison is described in section 5.2.3. The best-fitting block-wise model contained separate learning rates for win and loss outcome, one inverse temperature and a tendency parameter. The best-fitting constant inverse temperature model included separate learning rates for win and loss outcomes (per block) and separate inverse temperature parameters for wins and losses (constant across all task blocks).

6.2.7 Statistical analysis

We started by analysing the effect of tDCS on non-computational measures including total winnings, win-stay probability, no-loss-stay probability, loss-switch probability and no-win-switch probability. As a non-computational measure of positive bias, we included the proportion of win-driven choices. This measure was based on choices that participants made after a trial in which the win and loss occurred with the same shape. The proportion of "win-driven choices" is the number of trials in which participants chose the shape that was associated with both outcomes, divided by the total number of trials in which the win and loss were associated with the same shape.

Regarding computational parameters, learning rates and inverse temperature estimates were analysed. In addition, the adjustment of win and loss learning rates between volatile and stable blocks was calculated. *Win learning rate adjustment* was defined as the difference between the win learning rate in the wins-volatile block and the win learning rate in the losses-volatile block. *Loss learning rate adjustment* was defined as the difference between the loss learning rate in the losses-volatile block and the loss-learning rate in the wins-volatile block. To capture the extent to which learning rate adjustment was biased towards either win or loss outcomes, a measure for *learning rate adjustment bias* was calculated, which was defined as loss learning rate adjustment minus win learning rate adjustment.

All analyses were performed in RStudio (Version 1.4.1717, R 4.1.1). All dependent variables were analysed in mixed ANOVAs, with the main factor of interest being the factor tDCS Condition (real

vs. sham). To analyse the effect of tDCS during the stimulation period (or immediately afterwards for the "offline tDCS" condition), only data from the second and third task blocks were included in the analysis. For significant effects of tDCS, the same analysis was run in the second half of the task to test whether the effect outlasted the stimulation period. All analyses on non-computational measures included the between-subject factor Block Order (wins-volatile first vs. losses-volatile first) and the within-subject factors Volatility (wins-volatile and losses-volatile). To test our first hypothesis that online bifrontal tDCS would increase win learning rates, learning rate estimates were analysed in ANOVAs including all factors mentioned above and the additional within-subject factor Valence (win learning rate vs. loss learning rate). Significant two-way interactions of tDCS Condition with Valence or Volatility and three-way interactions between tDCS Condition, Valence and Volatility were followed up. To test whether tDCS had an effect on learning rate adjustment, learning rate adjustment was analysed in an ANOVA including factors mentioned above apart from Volatility. Significant interactions between tDCS Condition and Valence were followed up. To test whether tDCS had an effect on the learning rate adjustment bias, the learning rate adjustment bias was analysed in an ANOVA including all factors mentioned above apart from Volatility and Valence. Win-stay probability and no-loss-stay probability were analysed with Wilcoxon rank-sum tests due to violations of the normality assumption.

To test our second hypothesis, that the effect of online bifrontal tDCS would be cognitive-state dependent, we ran the same analyses in the "offline tDCS" condition. The effect of real vs. sham tDCS was compared between the two samples using t-tests where relevant.

All analyses were repeated after removing potential outliers. Statistics are reported for the entire dataset unless outlier removal had an impact on the results (figures including outliers are reported in the supplementary analyses section for all analyses outlier removal had an impact on). Since we were interested in the effect of tDCS, outliers were identified based on the difference between real and sham tDCS to ensure that datapoints which might bias the effect of tDCS were removed. A datapoint was identified as potential outlier if it was more than 1.5 times the interquartile range

below the first or above the third quartile. P-values below 0.05 are considered significant, whereas p-values between 0.05 and 0.10 are interpreted as trends.

Analyses of the questionnaire data collected before and after task performance are reported in the supplementary analysis section.

6.3 Results

6.3.1 Effects of online bifrontal tDCS

6.3.1.1 Non-computational measures

To test whether online bifrontal tDCS had an effect on non-computational measures, we analysed total winnings, stay and switch probabilities and proportion of win-driven choices in separate ANOVAs. There was no effect of tDCS on any non-computational measure (all $p > .05$).

6.3.1.2 Block-wise model

To test our main hypothesis that bifrontal tDCS would increase learning rates from wins, we conducted an ANOVA with tDCS Condition and Valence as factors of interest. After removal of five outliers (see Supplementary Figure 6.7), there was a significant main effect of tDCS ($F(1,32) = 5.7$, $p = .022$; before outlier removal: $F(1,37) = 1.5$, $p = .22$) but no interaction between tDCS Condition and Valence ($F(1,32) = 0.01$, $p = .90$). Real compared to sham tDCS increased learning rates independent of valence (Figure 6.2A).

To test whether tDCS might have an effect on learning rate adjustment, we analysed learning rate adjustment as well as the learning rate adjustment bias in separate ANOVAs. There was no main effect of tDCS on learning rate adjustment ($F(1,37) = 0.36$, $p = .54$) nor an interaction between

tDCS and Valence ($F(1,37) = 1.2$, $p = .26$)(Figure 6.2 B). There was also no effect of tDCS on the learning rate adjustment bias ($F(1,37) = 1.2$, $p = .26$)(Figure 6.2C).

We also analysed the inverse temperature to test whether tDCS had an effect on choice randomness: there was no effect of tDCS on the inverse temperature (all $p > .47$).

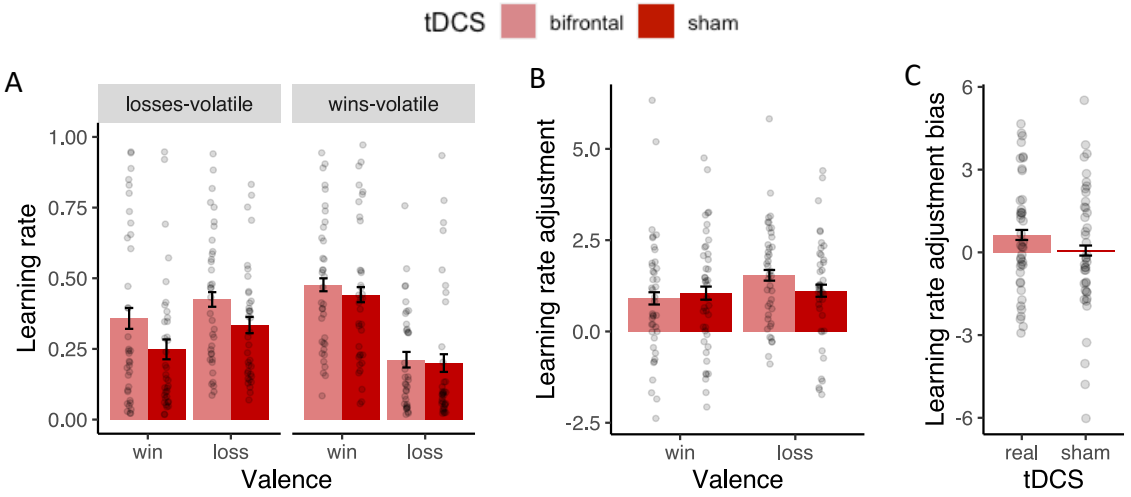

Figure 6.2. Effect of online tDCS on learning rates in the block-wise model. (A) Online bifrontal tDCS increased learning rates independent of valence. Online bifrontal tDCS had no effect on learning rate adjustment (B) or learning rate adjustment bias (C).

6.3.1.3 Constant inverse temperature model

To test our main hypothesis that tDCS would selectively increase the learning rate from wins, we ran an ANOVA with tDCS Condition and Valence as the main factors of interest. There was no significant main effect of tDCS nor an interaction between tDCS Condition and Valence, indicating that tDCS had no effect on learning rates (main effect of tDCS Condition: $F(1,37) = 0.3$, $p = .86$); tDCS Condition x Valence interaction: $F(1,37) = 0.26$, $p = .60$)(Figure 6.3A).

To test for a potential effect of tDCS on learning rate adjustment, we ran an ANOVA with learning rate adjustment as dependent variable. There was no main effect of tDCS Condition ($F(1,34) = 0.1$, $p = .70$), but there was a significant interaction effect between tDCS Condition and Valence

($F(1,34) = 7.47$, $p = .009$) after removal of three outliers (see Supplementary Figure 6.8)(before outlier removal: ($F(1,37) = 3.7$, $p = .059$). Real compared to sham tDCS led to an increase in loss learning rate adjustment ($F(1,34) = 6.0$, $p = .018$, *Cohen's d* = 0.46), and to a marginally significant decrease in win learning rate adjustment ($F(1,34) = 4.0$, $p = .051$, *Cohen's d* = -0.48) (Figure 6.3B). In line with this, real vs. sham tDCS had a significant effect on the learning rate adjustment bias ($F(1,34) = 7.4$, $p = .009$, *Cohen's d* = 0.65; before outlier removal: $F(1,37) = 3.7$, $p = .059$). Real vs. sham tDCS increased the learning rate adjustment bias, i.e. shifted learning rate adjustment bias from wins towards losses (Figure 6.3C). During sham tDCS, the learning rate adjustment bias was negative and significantly different from zero ($t(37) = -2.1$, $p = .037$), indicating that participants adjusted their loss learning rates significantly less than their win learning rate. During real tDCS, there was a trend towards a positive learning rate adjustment bias, i.e. towards higher loss than win learning rate adjustment ($t(37) = 1.8$, $p = .067$). To test whether the effect of tDCS on learning rate adjustment outlasted the stimulation period, we ran the same ANOVAs with learning rate adjustment (or learning rate adjustment bias) as dependent variable in the task blocks after the stimulation period (block 3 and 4). There was no significant interaction between tDCS Condition and Valence (or main effect of tDCS Condition)($F(1,34) = 1.3$, $p = 0.24$), indicating that the effect was confined to the period during stimulation (see Supplementary Figure 6.9).

As a control analysis, we also analysed the inverse temperature estimates which are unlikely to be affected by tDCS since they were estimated across the entire task. There was no effect of tDCS Condition on the inverse temperature (all $p > .29$).

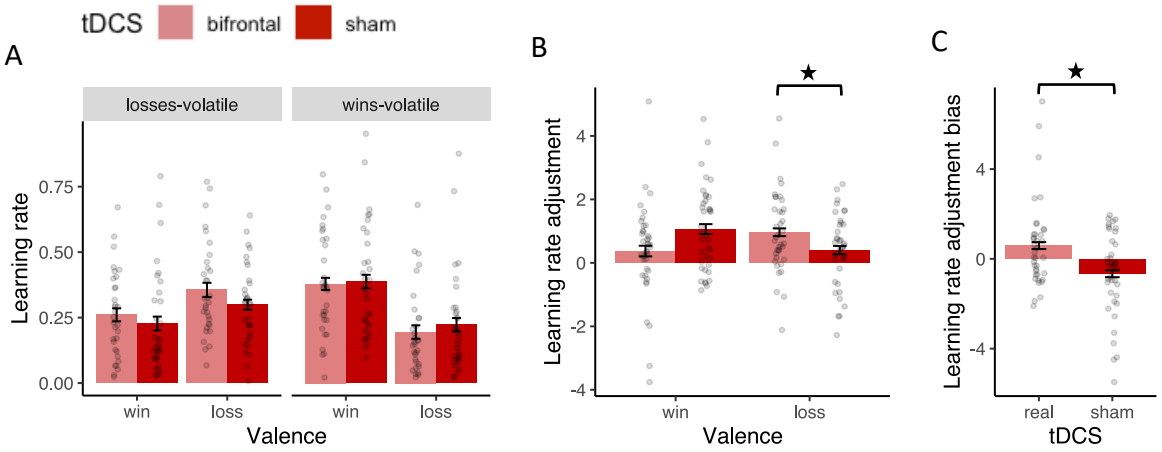

Figure 6.3. Effect of online bifrontal tDCS on learning rates in the constant inverse temperature model. (A) Online bifrontal tDCS had no effect on learning rates. (B) Real vs. sham tDCS significantly increased loss learning rate adjustment. TDCS marginally decreased win learning rate adjustment (p = .05). (B) Real vs. sham tDCS increased the learning rate adjustment bias, i.e. shifted learning rate adjustment away from wins towards losses.

6.3.2 Effects of offline bifrontal tDCS

We hypothesised that the effects of online bifrontal tDCS are specific to the stimulation time. We therefore conducted the same analyses in the group that received offline tDCS.

6.3.2.1 Non-computational measures

Similar to online tDCS, offline bifrontal tDCS had no effect on total winnings, win-stay probability, no-loss-stay probability, loss-switch probability or no-win-switch probability (all p > .48). Offline tDCS also had no effect on the proportion of win-driven choices (all p > .55).

6.3.2.2 Block-wise model

By contrast with online tDCS, offline tDCS had no significant effect on learning rates, learning rate adjustment, learning rate adjustment bias or inverse temperature (all p > .11)(Figure 6.4).

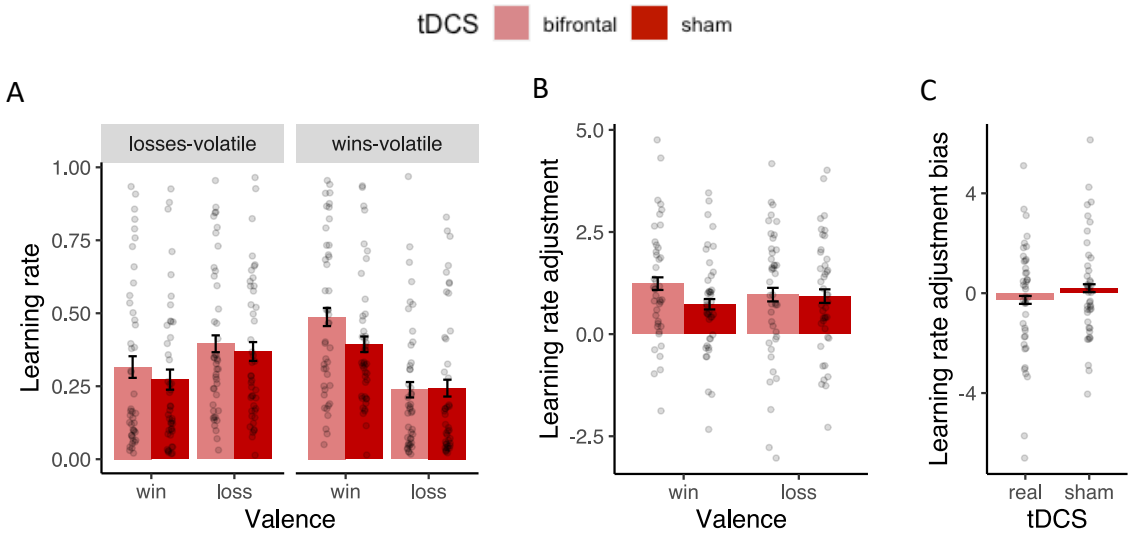

Figure 6.4. Effects of offline bifrontal tDCS on learning rates derived from the block-wise model. Real compared to sham tDCS had no effect on learning rates (A) learning rate adjustment (B) or learning rate adjustment bias (C).

6.3.2.3 Constant inverse temperature model

Offline tDCS had no significant effect on learning rates, learning rate adjustment, learning rate adjustment bias or inverse temperature (all $p > .16$)(Figure 6.5).

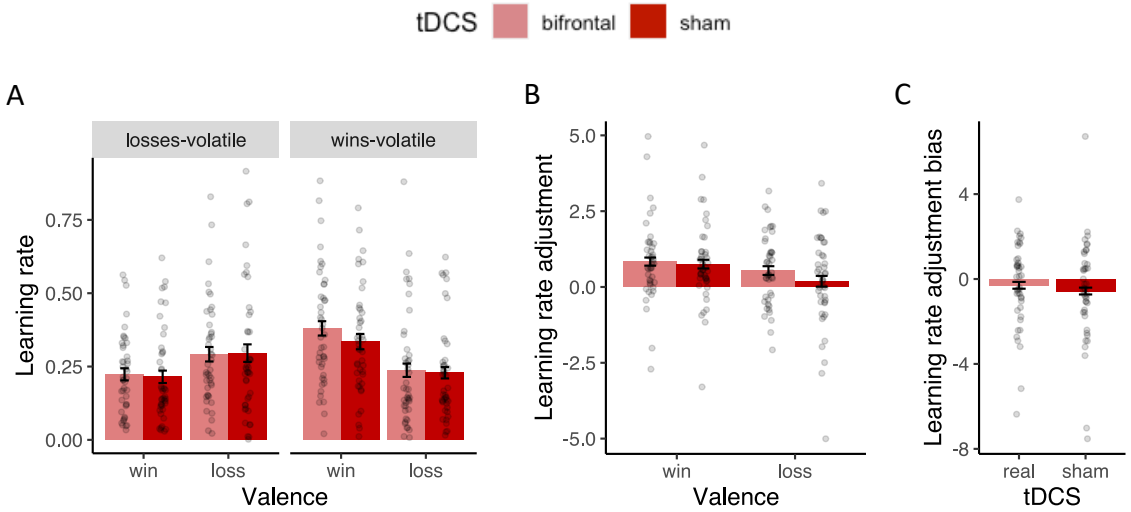

Figure 6.5. Effects of offline bifrontal tDCS on learning rates derived from the constant inverse temperature model. Real compared to sham tDCS did not have any significant effect on learning rates (A), learning rate adjustment (B) or learning rate adjustment bias (C).

The key effect of online tDCS was an increase in learning rate adjustment bias, an effect that was not significant in the offline condition. To test directly whether the effect of online bifrontal tDCS was specific to the stimulation time, we conducted a t-test to contrast the effect of real vs. sham tDCS on the learning rate adjustment bias between the online and offline tDCS conditions. The effect of online tDCS was marginally though not significantly larger than the effect of offline tDCS ($t(78.9) = -1.5$, $p = .063$ (Welch two sample t-test, one-sided))(Figure 6.6).

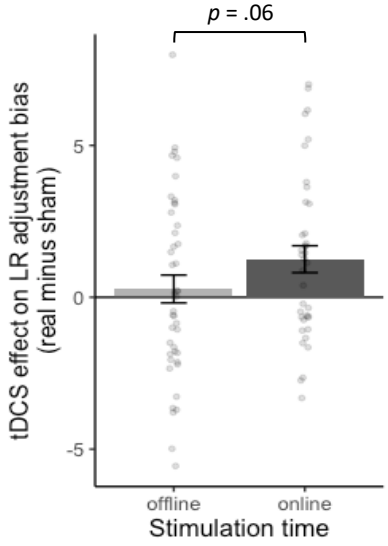

Figure 6.6. Comparison of the effects of online and offline bifrontal tDCS on learning rate adjustment bias. The effect of online tDCS on learning rate adjustment bias was marginally larger than the effect of offline tDCS.

6.4 Discussion

In this study, we investigated the effect of online bifrontal tDCS on reward learning in low mood. In contrast to our hypothesis, online bifrontal tDCS did not selectively increase reward learning rates. Instead, online bifrontal tDCS altered the adjustment of learning rates to changes in informativeness. Real vs. sham tDCS increased loss learning rate adjustment, thereby increasing the learning rate adjustment bias towards greater adjustment of learning rates from losses (relative to wins). This effect was specific to the cognitive state during stimulation since tDCS applied

before task performance had no such effect. However, this effect was only observed in one of the two modelling approaches that were used.

Negative biases are hypothesised to causally contribute to the development and maintenance of depressive symptoms (Disner et al., 2011). The reduction of negative biases is thought to be one mechanism underlying effective antidepressant treatments (Roiser, Elliott, & Sahakian, 2012) and is therefore a promising treatment target. In the previous chapter, we found that low mood was associated with altered learning rate adjustment. While healthy individuals adjusted win and loss learning rates to a similar extent, participants with low mood adjusted their loss learning rate to a smaller extent than their win learning rate and showed a shift away from loss towards win learning rate adjustment. The results from this study suggest that online bifrontal tDCS might normalise this behaviour. The decrease in loss learning rate adjustment in low mood present during sham tDCS was reversed during the application of real tDCS. During stimulation, there was actually a trend in the opposite direction, i.e. towards higher loss than win learning rate adjustment. The effect of tDCS on learning rate adjustment did not outlast the stimulation period. However, if stimulation during the task was repeated across multiple sessions, the effects of tDCS are likely to accumulate and persist.

These findings raise the question whether the normalisation of learning rate adjustment might be beneficial in depression treatment. Altered adjustment of learning rates to changes in informativeness has been suggested to be one potential mechanism underlying the development of negative biases (Browning et al., 2015; Gagne et al., 2020; Pulcu & Browning, 2017). More precisely, it has been suggested that negative biases might develop if individuals overestimate the informativeness of negative outcomes (Pulcu & Browning, 2017) and therefore show increased learning from negative outcomes (Aylward et al., 2019; Eshel & Roiser, 2010; Pike & Robinson, 2022). While we observed altered learning rate adjustment in individuals with low mood, learning from negative outcomes was not increased. Therefore, it is unclear if and how the changes in learning rate adjustment observed in low mood might lead to negative biases, and whether

normalising learning rate adjustment could reduce negative biases. Further research is needed to test for a causal relationship between learning rate adjustment and depressive symptoms. Future studies could train individuals with depressive symptoms in adjusting reward and punishment learning rates to informativeness, and test whether improved learning rate adjustment transfers to other tasks and might eventually improve mood.

TDCS affected learning rate adjustment only when applied *during*, but not when applied *before* task performance. This suggests that the cognitive state during the stimulation period is critical. In clinical trials, tDCS is usually applied at rest. However, tDCS has been shown to induce plasticity and learning effects (Fritsch et al., 2010; Reis et al., 2009), and might therefore be more effective if applied during a learning task relevant to depression. Based on the alterations in learning rate adjustment in low mood that we observed (Chapter 5), a suitable task might be a training paradigm which trains participants in appropriately adjusting behaviour to changes in informativeness of positive and negative outcomes. The results from the current study support the hypothesis that tDCS might increase the learning effect induced by such a training task. A future clinical trial could test whether bifrontal tDCS applied during task performance in this way is more effective than tDCS applied at rest (and tDCS applied before task performance).

Our findings raise the question why online bifrontal tDCS had a different effect in participants with low mood than in healthy participants. We speculate that this is the case because participants with low mood performed the task in a different way at baseline compared to healthy individuals. TDCS has been shown to interact with baseline behaviour. For instance, tDCS increased attention towards positive or negative stimuli, depending on the learning effect induced by an attentional bias modification task (Clarke et al., 2014). Participants with low mood adjusted their learning rates differently than healthy individuals, i.e. their baseline behaviour differed qualitatively from the behaviour of healthy individuals. The divergent effect of bifrontal tDCS in low mood might result from these differences in baseline behaviour. In individuals with low mood, bifrontal tDCS might normalise aberrant brain activity and thereby normalise information processing. The same tDCS

protocol applied to a healthy brain might increase a pre-existing optimism bias which is associated with intact mental health (Korn et al., 2014; Sharot et al., 2011).

As a limitation of the current findings, it should be noted that the effect of online bifrontal tDCS on learning rate adjustment was only observed in the model in which the inverse temperature parameters were kept constant across the entire task. This model should provide a robust estimate of the inverse temperature but is based on the assumption that the inverse temperature does not change between task conditions, which might be oversimplified (see section 5.4). In the block-wise model, we unexpectedly observed an increase in learning rates in response to real vs. sham tDCS. However, the effect only reached significance after removing a relatively large number of outliers (five) and was far away from reaching significance before outlier removal, which casts doubt on the reliability of this finding. Moreover, it should be noted that the reported p-values are not corrected for multiple comparisons. Our key finding, the increase in learning rate adjustment bias, would survive Bonferroni correction for 4 comparisons ($p = .009$, $p_{corr} = .036$; two hypotheses tested in two different models), whereas the increase in learning rates in the block-wise model would not survive this correction ($p = .022$, $p_{corr} = .088$). As discussed in Chapter 4, the addition of a separate dataset as control condition does not further increase the likelihood of these findings to be false-positive.

6.5 Conclusions

To conclude, we found that online bifrontal tDCS normalised altered learning rate adjustment in individuals with low mood. This effect was cognitive-state dependent since it was only present when tDCS was applied during (rather than before) task performance. The effects of combined tDCS and reinforcement learning might have potential clinical benefits beyond the application of tDCS at rest. Future clinical trials should test whether repeated sessions with combined tDCS and reinforcement learning induce changes in information processing which transfer to other tasks, and might ultimately improve mood.

6.6 Supplementary analyses

6.6.1 Additional questionnaires

Scores for the STAI state-anxiety scale and PANAS obtained before and after task performance are included in Table 6.2. To analyse potential effects of tDCS on mood, we entered the STAI state-anxiety score, the PANAS positive scale and PANAS negative scale into separate ANOVAs including the factors tDCS Condition (real vs. sham) and Time (before vs. after task performance). A significant interaction between Time and tDCS Condition would indicate that tDCS had an effect on mood. No significant interaction effect between Time and tDCS Condition were observed for any of the measures for online or offline tDCS (all $p > .05$).

Table 6.2. Mean (SD) for the STAI state anxiety scale and PANAS (low mood).

	Online tDCS		**Offline tDCS**	
	Before	After	Before	After
STAI-S				
Real	42.1 (10.0)	40.1 (8.3)	41.4 (8.2)	39.1 (9.1)
Sham	42.1 (11.0)	38.3 (9.1)	43.0 (11.6)	39.2 (10.7)
PANAS pos				
Real	24.2 (7.6)	21.5 (8.0)	24.2 (6.9)	23.7 (7.6)
Sham	24.1 (6.8)	22.3 (7.1)	23.0 (7.6)	22 (8.8)
PANAS neg				
Real	15.1 (5.9)	13.3 (4.1)	15.2 (3.9)	13.5 (3.9)
Sham	15.1 (5.8)	13.2 (4.1)	16.1 (6.1)	13,2 (4.9)

STAI-S: State-Trait Anxiety Inventory (State-anxiety scale, range: 20-80); PANAS: Positive and Negative Affect Scale (positive / negative affect scale, range: 10-50)

6.6.2 Outlier removal

Since we were primarily interested in the effect of real vs. sham tDCS, datapoints were identified as outliers if the difference between real and sham tDCS for a dependent variable was more than 1.5 times the interquartile range above the third, or below the first quartile. Outlier removal for learning

rates derived from the block-wise model and learning rate adjustment in the constant inverse temperature model is presented in Figure 6.7 and Figure 6.8, respectively.

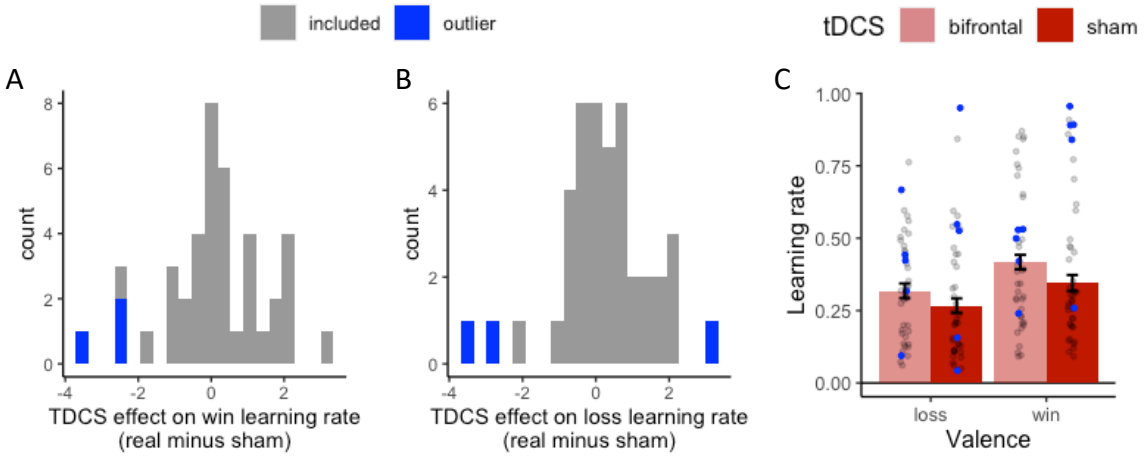

Figure 6.7. Outlier removal for the effect of tDCS on learning rates derived from the block-wise model (five outliers). Outliers were identified based on the difference in win (A) and loss learning rates (B) between real and sham tDCS. (C) shows the effect of tDCS on learning rates reported in the main text. Outliers are overlaid in blue.

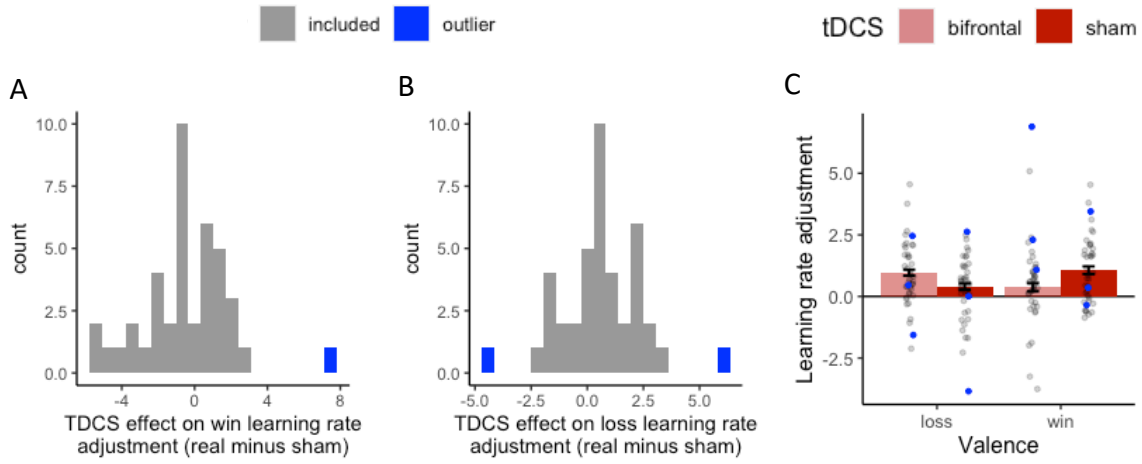

Figure 6.8. Outlier removal for learning rate adjustment based on the constant inverse temperature model (three outliers). Outliers were identified based on the difference in win (A) and loss learning rate adjustment (B) between real and sham tDCS. (C) shows the effect of tDCS on learning rate adjustment reported in the main text. Outliers are overlaid in blue.

6.6.3 tDCS effect after the stimulation period

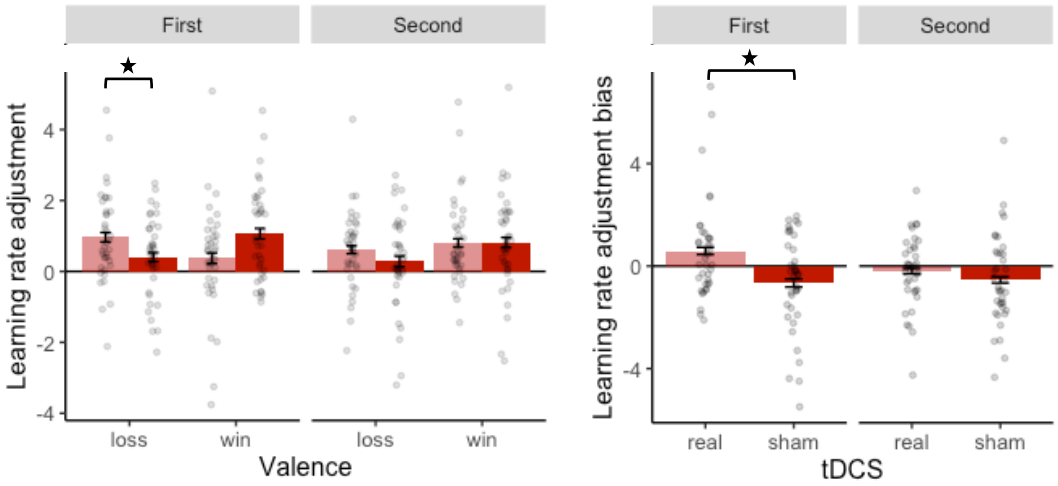

Figure 6.9. Effect of online bifrontal tDCS after the stimulation period. The effect of online bifrontal tDCS on learning rate adjustment (A) and learning rate adjustment bias (B) was only present during the stimulation time (First) but not after the stimulation had finished (Second).

7 Linking anxiety and anhedonia to cognitive task measures

Patients diagnosed with major depression show large heterogeneity in treatment response. Since patients also differ in their symptom profiles, different phenotypes of depression might exist. Patients with the same phenotype might share similar underlying psychopathology and might respond to the same treatment. Identifying each individual's phenotype could be helpful for guiding treatment selection, for instance to identify the optimal TMS target region. Since cognition plays an important role in depression, we propose to define phenotypes based on cognitive task performance. Parameters of a variety of cognitive tasks have been shown to correlate with symptoms of depression. However, it is unclear whether different task parameters capture distinct cognitive processes relevant to depression. The goal of this study was therefore to test whether parameters extracted from five different tasks differentially relate to different dimensions of clinical symptoms. Previous research indicates that the relative severity in symptoms of anxiety vs. anhedonia might have predictive potential for treatment response. Participants in an exploratory dataset (n = 402) collected online were therefore screened and recruited to ensure a sample that would show a wide range of relative severity of anxiety and anhedonia. Using factor analysis, we extracted two symptom factors which were positively correlated. The first factor was related to symptoms of anxiety and depression, predicted lower performance in reward learning (Probabilistic Instrumental Learning Task), less reliance on volatile loss probabilities (Volatility Task) and lower positive bias in emotion recognition (Facial Expression Recognition Task). The second factor was related to anhedonia and apathy, and predicted stronger reliance on volatile loss probabilities (Volatility Task) and better performance in punishment learning (Volatility Task, Probabilistic Instrumental Learning Task). Comparing the correlation profiles between the two symptom factors, our preliminary analyses suggest that the cognitive task that seemed to best differentiate between the two symptom factors was punishment learning in volatile and stable conditions. A confirmatory dataset will be collected to test whether these findings replicate. Future research can test for the existence of cognitive phenotypes (i.e. clusters based on task parameters) and their predictive value for treatment response.

7.1 Introduction

One of the biggest challenges in psychiatry is the heterogeneity in treatment response. Patients suffering from depression, for example, often need to undergo several treatment attempts before an effective treatment is found (Rush et al., 2006; Voineskos et al., 2020). To date, there is no effective system of determining who is likely to respond to which treatment.

Depression is characterised by a wide range of different symptoms. Patients with the same diagnosis ("depression") might have relatively little overlap in their symptoms. Given the heterogeneity in treatment response, this might suggest that different phenotypes with different underlying psychopathology exist (Beijers et al., 2019). Patients with the same phenotype of depression might share similar psychopathology and might respond to the same treatment. Identifying individuals' phenotype of depression could therefore help to guide treatment selection. While the left DLPFC is stimulated in the vast majority of TMS trials, recent research indicates that patients might differ in the TMS target region they best respond to (Downar et al., 2014; Drysdale et al., 2017). Phenotyping could help to identify the optimal target region for individual patients (Downar et al., 2014; Drysdale et al., 2017; Siddiqi et al., 2020b).

Previous studies have tried to define phenotypes purely on clinical symptoms which has not been found to predict response to antidepressant medication (Arnow et al., 2015; Uher et al., 2011). More recent research has focused on resting-state functional connectivity (Drysdale et al., 2017). However, fMRI scans are very unlikely to be included in standard psychiatric diagnostics. The current project is motivated by the goal of identifying depression phenotypes based on cognitive profiles derived from the performance of cognitive tasks (Baller et al., 2020). Since cognitive tasks can be performed online, cognitive phenotyping could easily be integrated into psychiatric diagnostics at low cost.

A growing body of evidence suggests that cognitive processing is altered in depression and anxiety (see (Bishop & Gagne, 2018) for a review). There is evidence for alterations in the processing of rewards and punishments (Aylward et al., 2019; Browning et al., 2015; Eshel & Roiser, 2010; Gagne, Agai, Ramiro, Dayan, & Bishop, 2022; Gagne et al., 2020; Harle, Guo, Zhang, Paulus, & Yu, 2017; Huys et al., 2013; Korn et al., 2014; Pike & Robinson, 2022; Pizzagalli et al., 2005), temporal discounting (Pulcu et al., 2014), meta-cognition (Rouault et al., 2018), and optimism bias (Korn et al., 2014). Importantly, these alterations in cognitive processing are hypothesised to causally contribute to the development and maintenance of depressive symptoms (Disner et al., 2011; Pulcu & Browning, 2017, 2019). In line with this, there is evidence that restoring these alterations, especially a reduction in negative biases, might be one mechanism underlying effective antidepressant treatments (Browning et al., 2019; Godlewska et al., 2016).

Previous research suggests that parameters in numerous experimental paradigms correlate with psychiatric symptoms. However, it is unclear to what extent these different task parameters capture distinct processes affected in depression. The aim of this project was therefore to test whether a selection of five different tasks that have previously been shown to capture cognitive processes affected in depression relate to different symptom dimensions of depression. Our approach builds upon previous work showing that alterations in cognitive processing relate to symptom dimensions cutting across boundaries of conventional psychiatric categories (Gillan et al., 2016; Rouault et al., 2018; Scholl et al., 2022).

Prior research on personalising non-invasive brain stimulation in the treatment of depression suggests that different subtypes of depression might differ in their relative severity of anxiety vs. anhedonia symptoms (Downar et al., 2014; Drysdale et al., 2017; Siddiqi et al., 2020b). The goal of this study was to identify cognitive task parameters which differentially relate to anxiety versus anhedonia. To maximise statistical power, we recruited a large sample of participants on an online platform and selected participants in such a way as to maximise the variance in relative severity of

anxiety and anhedonia symptoms (i.e. maximise the variance of the bivariate distribution of anxiety and anhedonia).

We have chosen two tasks which are well-established in psychiatric research. The *Probabilistic Instrumental Learning Task* (Pessiglione et al., 2006) assesses reward and punishment learning. Depression has been associated with lower performance (especially in reward learning) in this task (Kumar et al., 2018; Walsh, Browning, et al., 2018) and antidepressants have been found to normalise this deficit (Walsh, Browning, et al., 2018). The *Facial Expression Recognition Task* (Young et al., 1997) measures the processing of positive and negative facial expressions. Depression is associated with a response bias towards negative vs. positive facial expressions (Gur et al., 1992; Walsh, Browning, et al., 2018; Walsh, Huneke, et al., 2018), and antidepressants have been found to reduce this bias (Bhagwagar et al., 2004; Walsh, Browning, et al., 2018). Although both tasks have been applied in numerous studies, it is unclear which symptom dimension (anxiety or anhedonia) they would be most sensitive to.

In addition, we have chosen three more recently developed tasks which have been found to capture cognitive processes particularly relevant to either anhedonia or anxiety. The *Volatility Task* (Browning et al., 2015) captures how individuals adjust their learning behaviour to changes in the informativeness of punishments. Deficits in this ability have been related to symptoms of anxiety specifically (Browning et al., 2015), or symptoms of internalising symptomatology (i.e. symptoms of depression and anxiety) in general (Gagne et al., 2020). Deficits in adjusting learning to changes in informativeness is hypothesised to be a potential mechanism leading to negative biases (Browning et al., 2015; Pulcu & Browning, 2017, 2019). The *Wheel of Fortune Task* (Kolling et al., 2018) has been developed to investigate sequential decision making. This task requires participants to choose between accepting the current offer or keeping on searching the alternatives which might lead to a better offer. A recent study found that apathy was related to a specific parameter capturing decision inertia (inflexibility and repetitiveness in decision making), i.e. continuing to search the alternatives for longer than appropriate given the value of searching

(Scholl et al., 2022). The *Foraging Fish Game* is a paradigm currently under development which captures threat avoidance, and which is hypothesised to relate specifically to symptoms of anxiety.

The aim of this chapter was to explore relationships between psychiatric symptoms and various parameters derived from these tasks (except for the Foraging Fish Game, which is still under development). Four main hypotheses related to the key parameter for each task were tested: (1) Stronger symptoms of apathy/anhedonia would predict stronger decision inertia in the Wheel of Fortune Task; (2) Stronger symptoms of anxiety would predict lower learning rate adjustment between volatile and stable conditions in the Volatility Task; (3) Stronger internalising symptomatology (no specific hypothesis for anhedonia vs. anxiety) would predict lower performance in reward learning in the Probabilistic Instrumental Learning Task; (4) Stronger internalising symptomatology would predict a lower positive bias in the recognition of emotional face expressions in the Facial Expression Recognition Task.

7.2 Methods

7.2.1 Study design

This study was designed to include two samples, an exploratory sample and a confirmation sample. The exploratory sample will be used to explore relationships between psychiatric symptoms and task measures, test different analysis approaches, and generate specific hypotheses which will be pre-registered. The validation sample (not yet collected) will then be used to test these pre-registered hypotheses in an independent dataset. This chapter reports preliminary findings from exploring relationships between psychiatric symptoms and task measures in the exploratory dataset. The validation sample will be collected at later stages of the project.

7.2.2 Recruitment procedure

The study was conducted online via the Prolific recruitment platform (https://www.prolific.co/) between May and August 2022. Questionnaires were implemented using Qualtrics software (Qualtrics, Provo, UT, https://www.qualtrics.com), tasks were implemented on Pavlovia (https://pavlovia.org/). We aimed at recruiting a sample of 400 participants. To maximise the variance of anxiety and anhedonia symptoms, a screening was conducted including the State-Trait Inventory for Cognitive and Somatic Anxiety (STICSA)(Gros et al., 2007), Apathy Motivation Index (AMI)(Ang et al., 2017) and Temporal Experience of Pleasure Scale (TEPS)(Gard et al., 2006). These questionnaires were used to calculate two scores for each participant, one for anxiety (STICSA) and one for anhedonia (average of AMI and TEPS). Anxiety and anhedonia scores were transformed to values between 0 and 1. The scores were divided into three categories, low (0-0.25), medium (>0.25-0.5) and high (>0.5-1). The aim was to recruit 45 participants for each combination of categories for anxiety and anhedonia. Since some category combinations were very rare (especially low anhedonia – medium anxiety, and low anhedonia – high anxiety) the maximum number of participants per category combination was increased to 65 during the recruitment process (see Figure 7.1 for distribution of screening scores).

Participants who were selected based on their screening scores were invited to three testing sessions on three consecutive days (see Table 7.1 for an overview of the whole study procedure). Session 1 took place on the day of the screening. All three testing sessions started with the completion of further questionnaires (see Table 7.1) followed by the performance of one or two experimental tasks. Session 1 included the Wheel of Fortune Task and Volatility Task, session 2 the Probabilistic Instrumental Learning Task and Facial Expression Recognition Task, and session 3 the Foraging Fish Game. The order of tasks within session 1 and session 2 was randomised. Participants were informed that the three testing sessions would take place on three consecutive days. In case they were unable to complete the sessions on three consecutive days they were given up to seven days to complete the entire study.

To ensure participants were paying attention to the questionnaires and tasks, several quality checks were included throughout all phases of the study (see next section). Participants were only invited to the next session in the study procedure if they had passed all quality checks in the previous session. Participants who failed one or more of the quality checks were not invited to the following session but were still reimbursed for the parts they had already completed. Participants were paid £7.50 per hour. As an incentive to complete the entire study and perform well, participants were told that they would receive a bonus payment of up to £5 depending on task performance after completion of all three testing sessions. In reality, all participants who completed the study were paid a bonus of £5 independent of their task performance.

Table 7.1. Overview of questionnaires and tasks.

Questionnaires	Tasks
Screening	
State-Trait Inventory of Cognitive and Somatic Anxiety (STICSA)[1]	
Temporal Experience of Pleasure Scale (TEPS)[2]	
Apathy Motivation Index (AMI)[3]	
Testing Session 1 (1h)	
Demographics questionnaire	Volatility Task
Beck Depression Inventory (BDI-II)[4]	Wheel of Fortune Task
Quick Inventory of Depressive Symptomatology (QIDS (6 items))[5]	
Work and Social Adjustment Scale (WSAS)[6]	
Testing Session 2 (45 mins)	
Obsessive-Compulsive Inventory Revised (OCI-R)[7]	Facial Expression Recognition Task (FERT)
Short Scales for Measuring Schizotypy (SSMS; unusual experiences)[8]	Probabilistic Instrumental Learning Task (PILT)
Testing Session 3 (1h)	
Snaith-Hamilton Pleasure Scale (SHAPS)[9]	Foraging Fish Game
Experiences Questionnaire, Decentering Factor (EQ-D)[10]	
Intolerance of Uncertainty Scale (IUS)[11]	

1 (Gros, Antony, Simms, & McCabe, 2007)
2 (Gard, Gard, Kring, & John, 2006)
3 (Ang, Lockwood, Apps, Muhammed, & Husain, 2017)
4 (Beck, Steer, & Brown, 1996)
5 (Rush et al., 2003)
6 (Mundt, Marks, Shear, & Greist, 2002)
7 (Foa et al., 2002)
8 (Mason, Linney, & Claridge, 2005)
9 (Snaith et al., 1995)
10 (Fresco et al., 2007)
11 (Carleton, Norton, & Asmundson, 2007)

7.2.3 Sample

The screening was available to all participants on Prolific who had indicated that they were at least 18 years of age, resident in the UK and that their first language was English. Participants were screened in nine batches of up to 900 participants. Overall, 3829 participants completed the screening. 856 (22%) of these were invited to session 1 based on their screening scores. 689 (80%) of these participants completed session 1. 180 (26%) of these were excluded from the rest of the study due to failing at least one of the quality checks in session 1. 509 participants (74%) were invited to take part in session 2, 494 (97%) of whom completed session 2. 74 (15%) of these participants were excluded due to failing at least one quality check in session 2. 420 (85%) participants were invited to session 3, 402 (96%) of whom completed session 3. The final sample therefore consisted of 402 participants (which corresponds to 10% of participants who took part in the screening, and 58% of participants who completed session 1).

Although the distribution across anhedonia and anxiety categories was not uniform, the sample showed high variance in the bivariate distribution of anxiety and anhedonia scores (Figure 7.1A and 7.1B). The sample had a wide range of age (18 – 79 years, mean = 40 ± 13 years) and education levels (Figure 7.1C and 7.1E). 61% of participants were female (Figure 7.1F). 45% of participants reported a current or past diagnosis of a psychiatric disorder indicating that a substantial proportion of participants was characterised by clinically relevant symptom levels (Figure 7.1D).

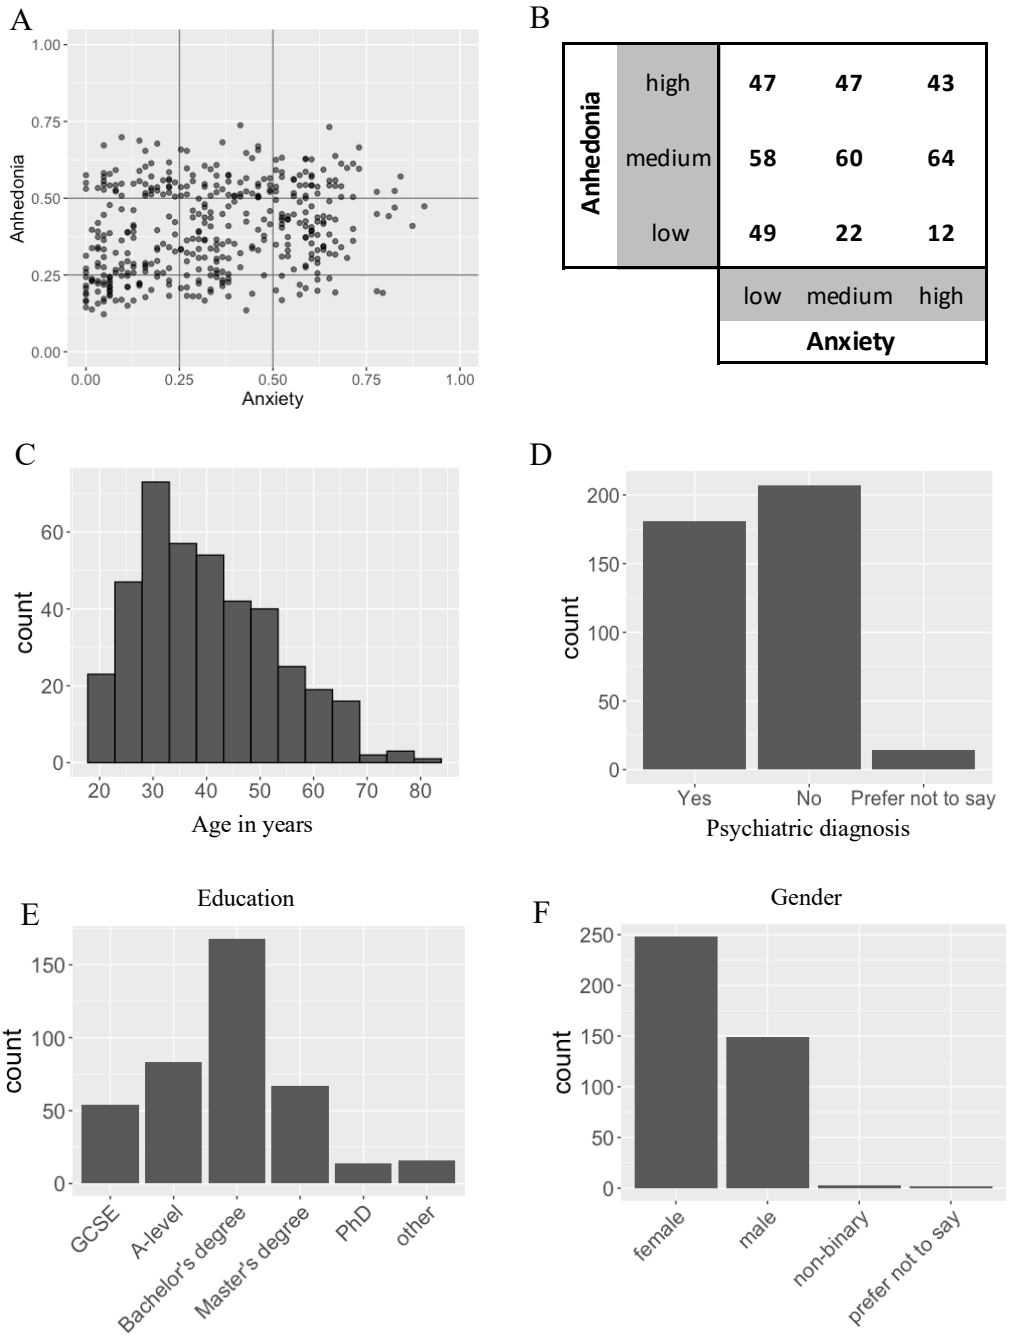

Figure 7.1. Demographic and clinical characteristics of the final exploratory sample (n = 402). (A) Distribution of anhedonia and anxiety symptoms assessed in the screening. (B) Number of participants per combination of anxiety and anhedonia symptom level. Distribution of (C) age, (E) education and (F) gender. (D) Proportion of participants who reported a current or past psychiatric diagnosis.

7.2.4 Experimental paradigms

The study included the performance of five tasks distributed over three testing sessions. All task measures included in the analyses are summarised in Table 7.2.

7.2.4.1 Wheel of Fortune Task

The Wheel of Fortune Task has been described in detail in (Kolling et al., 2018). Here, a shorter version with 100 trials was used. Participants were randomised to perform one out of two different task protocols. On each trial, a "wheel of fortune" was presented on the screen which was divided into different segments. The size of each segment represented the probability of the segment to be drawn, and numbers indicated the value of the segments. Participants were asked to choose between spinning the wheel of fortune, which resulted in a new offer, or banking the current offer. In some trials, spinning the wheel was associated with a cost (the amount varied across trials) which was displayed to the participant. The maximum number of available spins was also shown to the participant. The aim of the task was to maximise the number of points won.

The utility of spinning the wheel (or "searching") depended on several task parameters. This includes the *prospective value*, i.e. the value which stems from the opportunity of spinning the wheel multiple times. The prospective value for each trial was derived from a decision tree model taking into account the values and probabilities of the alternative options (Kolling et al., 2018). The prospective value might decrease over time as the number of available searches decreases (*prospective value difference since search 1*). The utility of spinning the wheel also depends on the current *offer value*, the *myopic value* (average value of the alternatives displayed on the wheel of fortune) and the *cost* of spinning the wheel. Another parameter which has been found to influence participants' likelihood of spinning the wheel is *number of previous searches*, which capture decision inertia (inflexibility and repetitiveness in decision making). The higher the number of previous searches already performed in one trial, the more likely participants tend to be to spin the

wheel again, which can be interpreted as a "stuck in a rut" bias. These parameters were included in a linear regression model to calculate the utility of spinning the wheel. The utility was then translated into a choice probability using a Softmax function including an inverse temperature parameter capturing choice randomness. This model was fitted to each participant's choices in the task using Stan. The model was fitted separately to initial searches (i.e. the first decision in each trial) and later searches (all choices in each trial after the initial decision). The regression weights indicate how much each participant's choices were influenced by each of the parameters. A previous study (Scholl et al., 2022) found that apathy and anhedonia were linked to different parameters included in the model for later searches, therefore only later searches were included in the analyses reported in this chapter. The study found that apathy was associated with one specific parameter in this task, the number of previous searches (*PrevSearch*)(Scholl et al., 2022). The higher the number of previous searches on one trial, the more likely participants were to search again, which corresponds to decision inertia (or a "stuck in a rut" bias). Individuals with higher apathy tended to be more influenced by the number of previous searches, i.e. showed stronger decision inertia. Based on this study, we hypothesised that higher scores of apathy (or anhedonia) would predict a stronger influence of the number of previous searches on participants' decisions. The relationship between all other model parameters and psychiatric symptoms was explored as well.

7.2.4.2 Volatility Task

The Volatility Task was a modified version of the paradigm used in (Browning et al., 2015). Participants were randomised to perform one out of two protocols including two blocks of 90 trials each. Participants started off with a total of 5,000 points and were instructed to try to avoid losing points throughout the task. On each trial, a green and a blue square were presented on the screen. A number inside the squares indicated which number of points might be lost if this option was chosen. After participants indicated their choice, a square appeared in the centre of the screen indicating the colour of the square that led to a loss on this trial. If the participant chose the losing

shape, the number associated with the chosen square was subtracted from the total. If the participant chose the other shape the total remained unchanged. The squares were associated with reciprocal loss probabilities which participants were supposed to learn over the course of the trials. In the "stable" condition, the square that lost more often remained the same across the whole block. In the "volatile" condition, the loss probability reversed a few times over the course of the block. The order of conditions was randomised across participants.

Participants' trial-by-trial choices were analysed with a computational model including a Rescorla-Wagner learning rule and Softmax function as described in (Browning et al., 2015). The model assumes that participants track the probability of the options to lead to a loss by updating their probability estimate, $r_{(i)}$, with a fraction of the prediction error ($r_{(i)} - out_{(i)}$) determined by the learning rate α:

$$r_{green\ square(i+1)} = r_{(i)} + \alpha(r_{(i)} - out_{(i)})$$

$R_{green\ square\ (i+1)}$ is the estimated probability on trial i+1 of the green square leading to a loss. The estimated probability of the blue square leading to a loss, $r_{blue\ square\ (i+1)}$ is (1- $r_{green\ square\ (i+1)}$). The estimated probabilities are combined with the loss magnitude displayed in the squares to calculate negative values of the two squares in the following way:

$$v_{green\ square\ (i+1)} = F(r_{(i+1)}, \gamma) m_{green\ square\ (i+1)}$$

$$v_{blue\ square\ (i+1)} = F(1 - r_{(i+1)}, \gamma) m_{blue\ square\ (i+1)}$$

Where $v_{\ green\ square\ (i+1)}$ and $v_{\ blue\ square\ (i+1)}$ correspond to the estimated negative values of the two squares on trial i+1, and $m_{green\ square\ (i+1)}$ and $m_{blue\ square\ (i+1)}$ correspond to the loss magnitudes displayed in the squares. F(r, γ) is a linear transformation with output values between 0 and 1:

$$F(r, \gamma) = \max\left[\min[(\gamma(r - 0.5) + 0.5), 1], 0\right]$$

The risk preference parameter allows the model to capture whether participants' choices are primarily influenced by the estimated outcome probability ($\gamma > 1$) or outcome magnitude ($\gamma < 1$). The estimated negative values are transformed into action probabilities using a Softmax function:

$$P_{(choice=green\ square)} = \frac{1}{1 + \exp\left(-\beta(v_{blue\ square} - v_{green\ square}\right)}$$

Where β is the inverse temperature which captures to what extent participants base their choices on the expected negative values.

This model was fitted separately to the volatile and stable condition by omitting the first 10 trials of each block. This resulted in three parameter estimates per participant per condition (α, β and γ) which were log-transformed. Relationships with symptom factors scores were assessed for each parameter in the volatile and stable condition as well as their average across volatile and stable conditions. Our main hypothesis was that higher anxiety would be associated with lower learning rate adjustment. Therefore, the difference in learning rate between volatile and stable conditions was derived (volatile minus stable).

7.2.4.3 Probabilistic Instrumental Learning Task (PILT)

The PILT was a modified version of the paradigm used in (Pessiglione et al., 2006). Participants were randomly assigned to one out of two protocols. The protocols included three blocks of 60 trials. On each trial, one out of two symbol pairs was presented on the screen. Participants were asked to choose one symbol. In win trials (when the "win pair" was presented) there was either a win outcome (+20 points) or no outcome (0 points). In loss trials (when the loss pair was presented), there was either a loss outcome (-20 points) or no outcome. Participants were supposed to learn over time which symbol in the win pair is more likely to win, and which symbol in the loss pair is less likely to lose. Win trials and loss trials were presented interleaved in random order (which was fixed for each protocol).

Performance in win and loss trials was measured as the percentage of correct choices (i.e. the shape that wins more often or loses less often) in the last 20 win trials and last 20 loss trials of each block. The first 10 trials of each block were omitted since this measure was intended to measure a stable performance level that was reached over the course of the trials. The main measure of interest was the performance in win trials. A positive bias measure was created by calculating the difference in this performance measure between win and loss trials (win minus loss trials). As a second measure of choice behaviour, choice consistency for win and loss trials was calculated as the percentage of trials in which the same shape as on the previous trial was chosen. A positive bias in choice consistency was defined as the log-transformed ratio of consistency in win vs. loss trials (consistency win / consistency loss). Reaction times for win and loss trials as well as a positive bias in reaction times (log(loss reaction times / win reaction times)) were included as additional measures.

7.2.4.4 Facial Expression Recognition Task (FERT)

The FERT has been adapted from (Young et al., 1997). Participants were randomly assigned to one out of two protocols. Each protocol included two blocks of 62 and two blocks of 63 trials. On each trial, a face with an emotional expression was presented for 500ms. Participants were asked to indicate the emotion of the face as quickly and accurately as possible by clicking onto the corresponding button. Emotions included happy, surprise, anger, fear, sadness and disgust. Ten neutral faces were included as well. Each emotion (apart from neutral) was presented at 10 different intensities, by four different actors (i.e. 4 x 10 images per emotion). The response buttons were presented underneath the image arranged in a circle around a fixation cross. Participants had to move the cursor back onto the fixation cross before the start of a new trial.

As accuracy measure, the unbiased hit rate (based on signal-detection theory (Heeger & Landy, 1997)) was calculated for each emotion in the following way:

$$Unbiased\ hit\ rate = percentage\ emotion\ correctly\ identified * \frac{number\ of\ trials\ emotion\ correctly\ chosen}{number\ of\ trials\ emotion\ chosen}$$

The unbiased hit rate for each emotion (apart from neutral) averaged across intensities was included in the analysis. The main measure of interest was a positive accuracy bias, defined as the log-transformed ratio of unbiased hit rate averaged across positive emotions (happy, surprise) divided by averaged hit rate for negative emotions (fear, sadness, anger, disgust). For each emotion, a mean reaction time was calculated by averaging across all trials in which the emotion was identified correctly. A positive reaction time bias was calculated as the log-transformed ratio of the mean reaction time for negative emotions divided by the mean reaction time for positive emotions.

7.2.4.5 Foraging Fish Game

In this task, participants played a fish harvesting algae in the ocean. They had to periodically check for the presence of predators approaching their fish and hide from them when they got close to them. This task is still under development and is therefore omitted from the remainder of this chapter. Preliminary analyses suggest that some task parameters specifically capture symptoms of anxiety.

7.2.5 Quality checks

Questionnaires. Two or three attention check items were included in the questionnaires for the screening and each testing session. These included items like "If you are still paying attention please choose 'moderately' here". Correct answers in all attention check items were required to pass the questionnaire quality check.

Each task started with written instructions and practice trials. Each task included multiple-choice questions to check participants' task understanding. If participants indicated an incorrect answer, the instructions were displayed again. The experiment only proceeded to the actual task protocol once all questions had been answered correctly. The questions were chosen in a way to cover all relevant aspects of the task instructions.

Wheel of Fortune Task. The task requires a certain number of choices for different choice categories to ensure the model can be fitted. Participants passed this criterion if they had at least 13 accept and spin choices in initial (first choice on each trial) as well as in later searches (choices after the first decision to spin in a trial). An additional quality check for task understanding was included. Participants passed this check if the acceptance rate for the 10 highest offer values was higher than the acceptance rate for the 10 lowest offer values.

Volatility Task. There are two strategies that participants might apply in this task, choosing the shape with the smaller loss magnitude and/or choosing the shape with smaller loss probability. Participants passed the quality check if, in each of the two conditions (stable or volatile), they applied at least one of these strategies in at least 55% of the trials.

PILT. Participants passed the quality check if their overall choice accuracy was at least 55%.

FERT. Participants passed the quality check if their average classification accuracy was at least 30% (chance level: 14%).

Table 7.2. Overview of task measures.

Task	Parameter	Description	Derived from	Interpretation	Main hypotheses
Wheel of Fortune Task	Offer	Value of the current offer which can be banked	Regression model predicting utility of searching	Negative predictor	
	Cost	cost of spinning the wheel (constant within each trial)		Negative predictor	
	Myopic value	Average value of the alternatives displayed on the wheel of fortune		Positive predictor	
	Prospective value search 1	Value of spinning the wheelon the first choice in each trial stemming from the opportunity of spinning the wheel multiple times; derived from decision tree model based on the values and probabilties of the alternatives		Positive predictor	
	Decrease in prospective value since search 1	The higher the number of searches participants have already performed on a trial, the more likely they are to search again ("stuck in a rut" bias)		Negative predictor	
	Number of previous searches	The higher the number of searches participants have already performed on a trial, the more likely they are to search again ("stuck in a rut" bias)		Positive predictor	Individuals with higher apathy should be more strongly influenced by the number of previosu searches.
Volatility Task	LR volatile	Learning rate in volatile condition	Rescorla-Wagner learning model with Softmax function		
	LR stable	Learning rate in stable condition			
	LR mean	learning rate mean across both conditions			
	LR adjustment	Difference in learning rate between volatile and stable condition (volatile minus stable)		Positive value: Learning rate in volatile higher than in stable condition	Individuals with stronger symptoms of anxiety should adjust their learning rates to a smaller extent.
	gamma volatile	Weight on probability vs magnitude in volatile condition		Higher value: more weight on probability	
	gamma stable	Weight on probability vs magnitude in stable condition		Higher value: more weight on probability	
	gamma mean	Weight on probability vs magnitude averaged across both conditions		Higher value: more weight on probability	
	beta volatile	Inverse temperature in volatile condition			
	beta stable	Inverse temperature in stable condition			
	beta mean	Inverse temperature averaged across both conditions			

Table 7.2 (continued). Overview of task measures.

Task	Parameter	Description	Derived from	Interpretation	Main hypotheses
Proabilistic Instrumental Learning Task	RT win trials	Reaction time in win trials	Raw data		
	RT loss trials	Reaction time in loss trials			
	Positive bias RT log	Positive bias in reaction times (Log-transformed ratio of reaction time in win vs loss trials)		Positive values indicate faster reactions in win vs loss trials	
	correct win last 20	Measure for performance in win trials (percentage of correct choices in the last 20 win trials in each block)			Stronger internalising symptomatology should be associated with lower performance in win trials.
	correct loss last 20	Measure for performance in loss trials (percentage of correct choices in the last 20 loss trials in each block)			
	Positive bias correct choices diff	Measure for bias towards better performance in win vs loss trials (correct_win_last_20 minus correct_loss_last_20)		Positive value: higher accuracy in win than in loss trials	
	consistency win	Choice consistency in loss trials (percentage of choices equal to choice on previous trials)			
	consistency loss	Choice consistency in win trials (percentage of choices equal to choice on previous trials)			
	consistency positive bias log	Positive bias in choice consistency (log-transformed ratio of conistency win / consistency loss)			
Emotional Face Recognition Task	UHR [emotion]	Unbiased hit rate for each emotion	Raw data		
	RT [emotion]	Reaction time average for each emotion (includes only trials in which emotion has been correclty identified)			
	UHR positive bias log	Positive bias in unbiased hit rate (log-transformed ratio of average unbiased hitrate for positive emotions / average unbiased hitrate for negative emotions)		Positive value: higher accuracy in positive than negative emotions	Stronger internalising symptomatology should be associated with a weaker positive accuracy bias.
	Positive bias RT log	Positive bias in reaction times (log-transformed ratio of average reaction time for negative emotions / average reaction time for negative emotions)		Positive value: faster reaction times for positive than negative emotions	

7.2.6 Analysis

All analyses were conducted in *RStudio* (version 1.4.1717, *R* version 4.1.1). The analysis consisted of two steps: factor analysis to extract symptom factors, and regression analyses to test for relationships between symptom factor scores and task measures. As a first step, exploratory factor analysis was run on the single items of the mood-related questionnaires, i.e. STICSA, TEPS, AMI and Beck Depression Inventory II (BDI)(Beck et al., 1996). Since this analysis largely recovered the subscales of these questionnaires, a second exploratory factor analysis was run on the questionnaire subscales to further reduce the dimensionality of the data. For both exploratory factor analyses, parallel analysis was used to determine the optimal number of factors (*fa.parallel* from *psych* package), and factor analysis with oblique rotation (*fa* from *psych* package, rotation "*oblimin*") was used to extract the factors. Absolute factor loadings of 0.4 or above are considered significant. Factor scores were calculated using regression-based weights.

Regressions were run using *brms* which implements Bayesian regression models using *Stan*. Regressions were run to predict task measures based on symptom scores, and to predict symptom scores based on task measures. All continuous variables were z-transformed. Age (continuous regressor), gender (categorical regressor) and education (ordinal regressor) were included in all regressions as control variables. To predict task measures, a regression analysis was run for each task measure separately, including the two symptom factors as regressors of interest, and the demographic variables as control regressors. Most analyses also included a regressor controlling for task performance since we were interested in the relationship between symptoms and specific cognitive processes rather than task performance in general (Wheel of Fortune Task: inverse temperature parameter; Volatility Task: mean inverse temperature across the two task blocks for the analyses on learning rates and gamma parameter, none for analyses on inverse temperature; PILT: none since task performance was the measure of interest; FERT: unbiased hit rate average across emotions for all unbiased hit rate measures; mean reaction time for all reaction time measures). To test the two hypotheses specific to anxiety vs. anhedonia symptoms (1. Apathy (or

anhedonia) predicts higher influence of the number of previous searches in the Wheel of Fortune Task, 2. Anxiety predicts lower learning rate adjustment in the Volatility Task), additional regressions were run including specific questionnaire subscales as regressors. For all task measures that were significantly predicted by either of the two symptom factors, correlations and partial correlations were run as control analyses.

To test whether task measures can predict symptom factor scores, separate regressions were run including one of four a priori selected main measures of interest as regressor to predict scores of either symptom factor (i.e. *number of previous searches* in the Wheel of Fortune Task, *learning rate adjustment* in the Volatility Task, *performance in win trials* in the PILT, *positive bias in unbiased hit rate* in the FERT).

For all regression analyses, the mean regression weight estimate and two-sided 95% credible interval are reported (mean, [lower boundary, upper boundary]). A regression weight is considered significant if the 95% credible interval excludes zero.

7.3 Results

7.3.1 Factor analysis of mood questionnaires

As a first step, factor analysis was run on the single items of all mood questionnaires, i.e. STICSA, TEPS, AMI and BDI. Parallel analysis suggested an optimal number of 8 factors. Further details including the factor loadings are included in the supplementary analysis section. The extracted factors largely recovered the subscales of the questionnaires, apart from the BDI (all BDI items loaded on the same factor).

To further reduce the dimensionality of the data, the subscales of the STICSA, TEPS and AMI, as well as the total BDI scores were entered into a second factor analysis (see Figure 7.2A for the

correlation structure). Parallel analysis revealed an optimal solution with two factors. Exploratory factor analysis was conducted with oblique rotation, allowing the factors to be correlated. The two factors together explained 63% of variance in the questionnaire subscales and were positively correlated (r = .29). Based on the observed loadings, we labelled the factors "Anxiety-Depression" and "Anhedonia-Apathy" (Figure 7.2B). The three AMI subscales loaded positively on the Anhedonia-Apathy factor. The TEPS subscales which are coded in the opposite direction (higher values indicate less severe symptoms) loaded negatively on the Anhedonia-Apathy factor. The two STICSA subscales and BDI score loaded positively on the Anxiety-Depression factor. Interestingly, the AMI "emotional sensitivity" subscale which loaded positively on the Anhedonia-Apathy factor, loaded negatively on the Anxiety-Depression factor. At first glance this might seem counter-intuitive since the two factors are positively correlated. Looking at the items of this subscale (e.g. "After making a decision, I will wonder if I have made the wrong choice" or "I feel sad or upset when I hear bad news"), a plausible interpretation of the sign of loadings would be that individuals scoring high on Anxiety-Depression tend to worry a lot about their own decisions and about others, while individuals scoring high on Anhedonia-Apathy are indifferent.

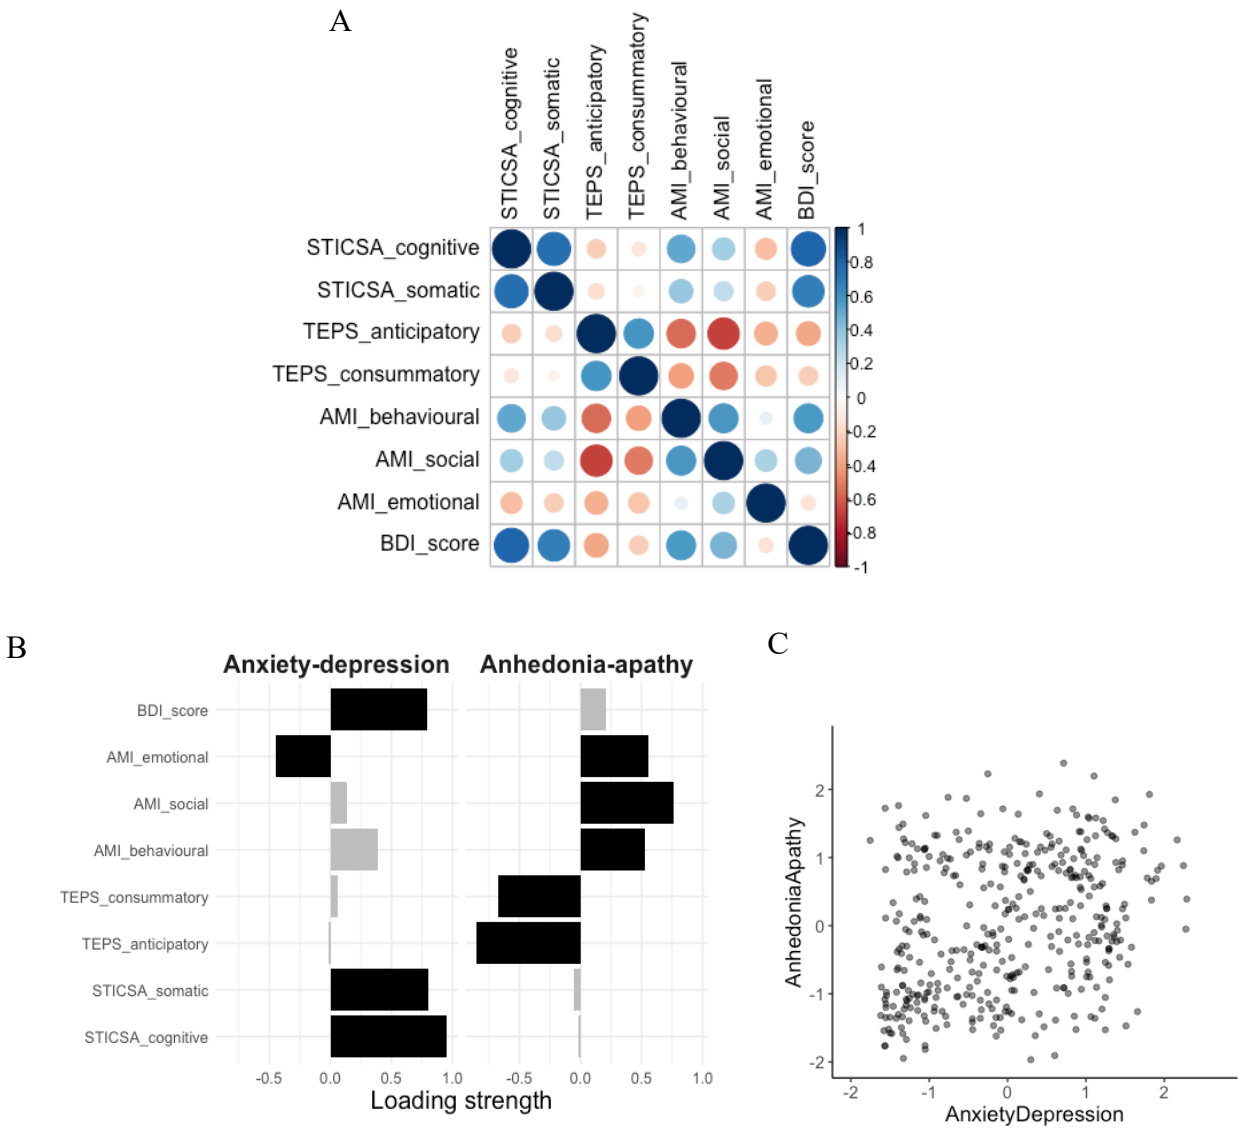

Figure 7.2. Exploratory factor analysis of the questionnaire subscales. (A) Correlation matrix of the questionnaire subscales.(B) Loadings of the questionnaire subscales. (C) Distribution of factor scores.

7.3.2 Predicting task measures based on symptom factor scores

To test for relationships between the two symptom factors and task measures, linear regressions were run to predict individual task measures based on the factor scores (Figure 7.2C). For each task measure a separate regression analysis was run using the Anxiety-Depression and Anhedonia-Apathy factors as regressors of interest.

7.3.2.1 Wheel of Fortune Task

For the Wheel of Fortune Task, we specifically predicted that individuals with higher apathy/anhedonia scores would be more likely to search, the higher their number of previous searches (*PrevSearch* parameter)(hypothesis 1). In contrast to our hypothesis, the Anhedonia-Apathy factor did not significantly predict this parameter (-0.06, [-0.17,0.05])(Figure 7.3). None of the TEPS and AMI subscales predicted the number of searches significantly. On a descriptive level the effect seems to be in the opposite direction, i.e. individuals with higher Anhedonia-Apathy scores tended to be less influenced by the number of previous searches than individuals with lower Anhedonia-Apathy scores. In line with this, Anhedonia-Apathy scores showed a trend towards predicting a conceptually related parameter, the *change in prospective value since search 1* ("*ProspValDiffSinceSearch1*")(-0.10, [-0.21,0.01]). A decrease in prospective value should prevent individuals from continuing to spin the wheel of fortune. On a trend-level, individuals with high Anhedonia-Apathy scores showed a more negative influence of decreases in prospective values, i.e. they were more likely to stop searching in response to decreases in prospective value than individuals with low Anhedonia-Apathy scores. This speaks against the hypothesis of anhedonia/apathy being associated with decision inertia.

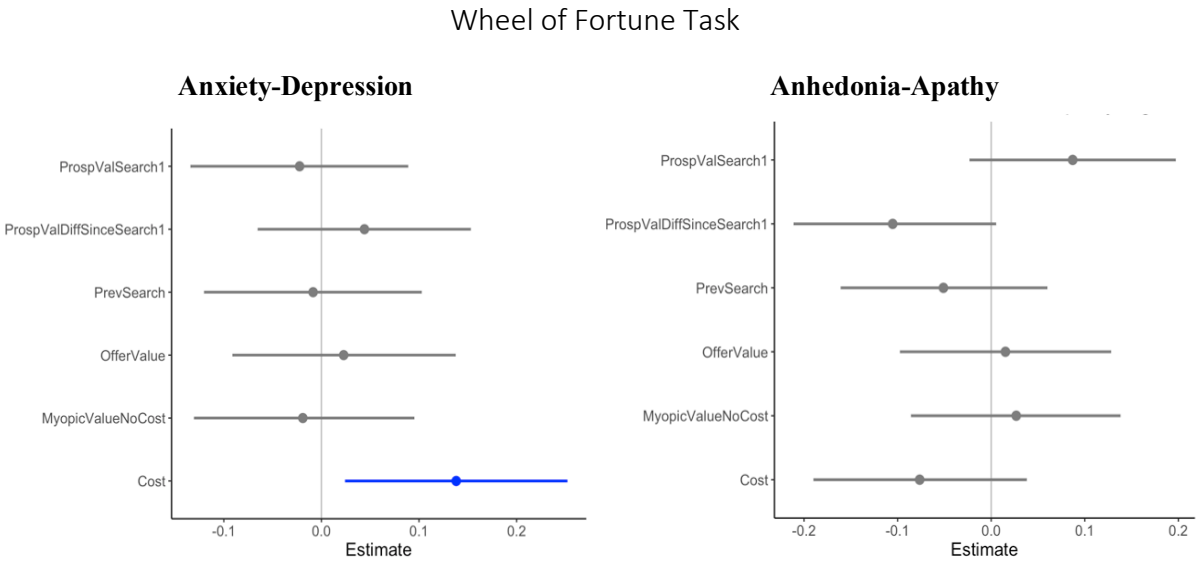

Figure 7.3. Regression weights for the Anxiety-Depression and Anhedonia-Apathy factors to predict each of the task measures extracted from the Wheel of Fortune Task. Error bars depict the 95% credible interval. Note that a more positive value for *Cost* corresponds to cost *in*sensitivity.

In an exploratory analysis, we found that the Anxiety-Depression factor (Figure 7.3) significantly predicted higher insensitivity to costs, i.e. individuals with higher Anxiety-Depression scores were less likely to be deterred by high costs of spinning the wheel than individuals with lower Anxiety-Depression scores (0.13, [0.01,0.24]). To test which questionnaire subscales contributed to predicting cost sensitivity, additional regressions were run separately for each of the subscales loading on Anxiety-Depression. The BDI score was the only significant predictor (0.12, [0.01,0.22]) indicating that the relationship between Anxiety-Depression scores and cost sensitivity was mainly driven by BDI scores. Further correlations were computed as control analyses. Without controlling for any variables, Anxiety-Depression scores showed a non-significant positive correlation with *Cost* (r = .07, p = .13)(Figure 7.4A). After regressing out the control variables age, gender and education, the residuals of *Cost* showed a trend towards a positive correlation with Anxiety-Depression (r = .09, p = .067)(Figure 7.4B). Since insensitivity to costs has previously been linked to higher compulsivity (Scholl et al., 2022), we tested for a correlation between Obsessive Compulsive Inventory (OCI)(Foa et al., 2002) scores and *Cost*. In line with previous work, the OCI score showed a trend towards a positive correlation with the Cost residuals (r = .08,

p = .11)(Figure 7.4C). After regressing out the demographic variables and OCI score from Cost, the strength of the correlation between Anxiety-Depression scores and Cost was reduced (r = .03, p = .47)(Figure 7.4D). The relationship between Cost and Anxiety-Depression scores was analysed in a second regression including an additional regressor to control for OCI score. Neither the weight for Anxiety-Depression (0.11, [-0.05, 0.26]) nor for OCI score (0.03, [-0.12, 0.17]) were significant, indicating that the relationship between Anxiety-Depression scores and Cost was driven by variability shared between these variables.

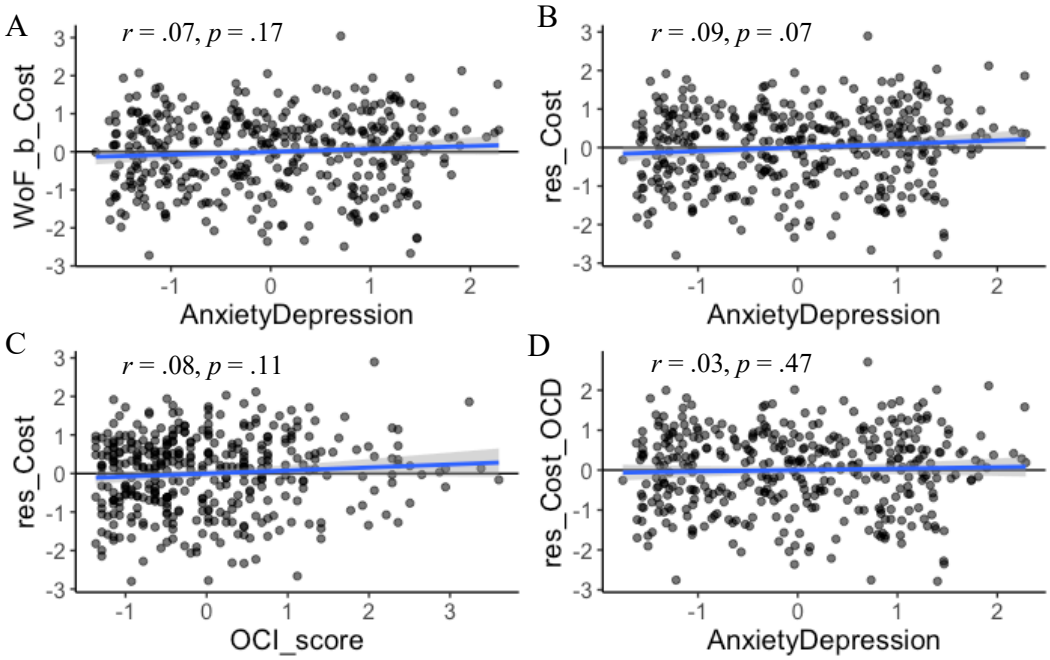

Figure 7.4. Partial correlations between Cost and Anxiety-Depression or OCI scores. (A) Correlation between Anxiety-Depression scores and Cost without controlling for any variable. (B) Partial correlation between Anxiety-Depressions scores and Cost after regressing out the demographic variables from Cost. (C) Partial correlation between OCI score and Cost after regressing out the demographic variables from Cost. (D) Partial correlation between Anxiety-Depression scores and Cost after regressing out demographic variables and OCI score from Cost.

7.3.2.2 Volatility Task

For the Volatility Task, we hypothesised that higher anxiety would predict lower learning rate adjustment (hypothesis 2). In contrast to our hypothesis, Anxiety-Depression did not significantly

predict learning rate adjustment, with a weight estimate around zero (0.01, [-0.09, 0.12])(Figure 7.5). The STICSA score, or either of the two subscales, which should represent a more specific measure of anxiety, did not significantly predict learning rate adjustment either.

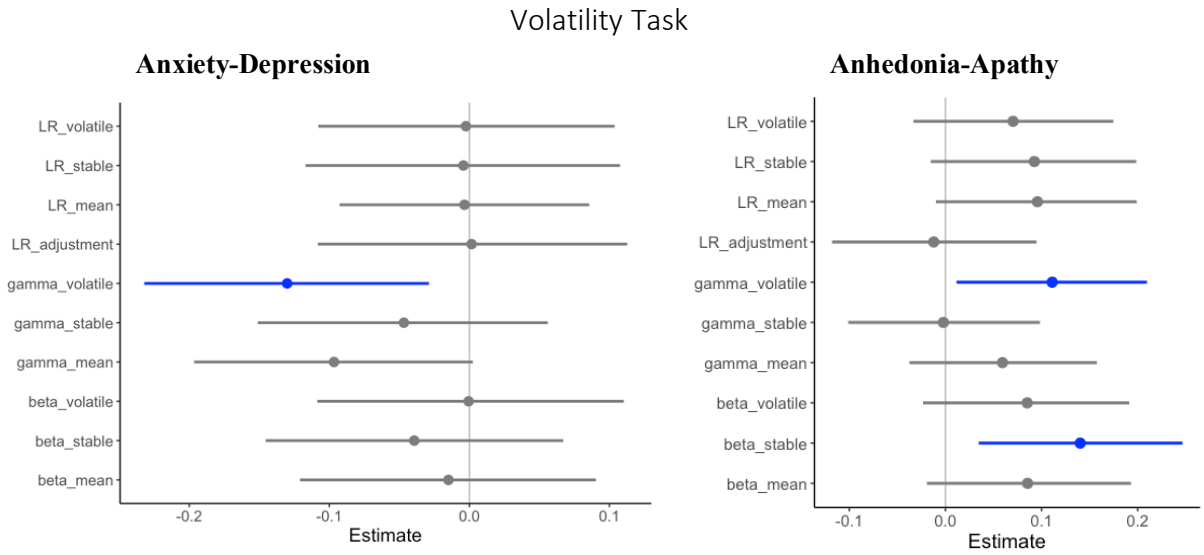

Figure 7.5. Regression weights for Anxiety-Depression and Anhedonia-Apathy for predicting individuals task measures extracted from the Volatility Task. Error bars indicate 95% credible intervals.

In an exploratory analysis, we found that Anxiety-Depression predicted a lower gamma parameter in the volatile condition (-0.12, [-0.22, -0.02]), i.e. individuals with high Anxiety-Depression scores put less weight on loss probability and more weight on loss amplitude in the volatile condition. Although this finding was unexpected, it is compatible with previous research which found that anxiety might be related to difficulties adjusting learning to volatile environments (Browning et al., 2015; Gagne et al., 2020). Individuals with high Anxiety-Depression scores might have difficulties tracking the loss associations in the volatile condition and might therefore base their choices primarily on loss magnitude. In contrast, Anhedonia-Apathy predicted a higher gamma parameter in the volatile block, i.e. more weight on loss *probability* (0.11, [0.01, 0.20]). Without controlling for any variables, *Gamma-volatile* showed a non-significant negative correlation with Anxiety-Depression scores (r = -.07, p = .14)(Figure 7.6A), and a non-significant

positive correlation with Anhedonia-Apathy scores (r = .02, p = .65)(Figure 7.6B). After controlling for demographic variables, the strength of these correlations increased (Anxiety-Depression: r = -.09, p = .06, Figure 7.6C; Anhedonia-Apathy: r = .07, p = .15, Figure 7.6D).

To test which subscales contributed to predicting *Gamma-volatile*, regressions were run for each subscale separately. No questionnaire subscale loading on the Anxiety-Depression factor significantly predicted *Gamma-volatile*. The regression weight for the BDI score (-0.09, [-0.18, 0.002]) and STICSA cognitive subscale (-0.09, [-0.18, 0.001]) were closest to significance and, in line with the findings above, both predicted a lower *Gamma-volatile*. AMI social was the only subscale loading on the Anhedonia-Apathy factor that significantly predicted *Gamma-volatile* (0.13, [0.03, 0.22]). In line with the findings above, higher scores on the AMI social subscale predicted higher *Gamma-volatile*.

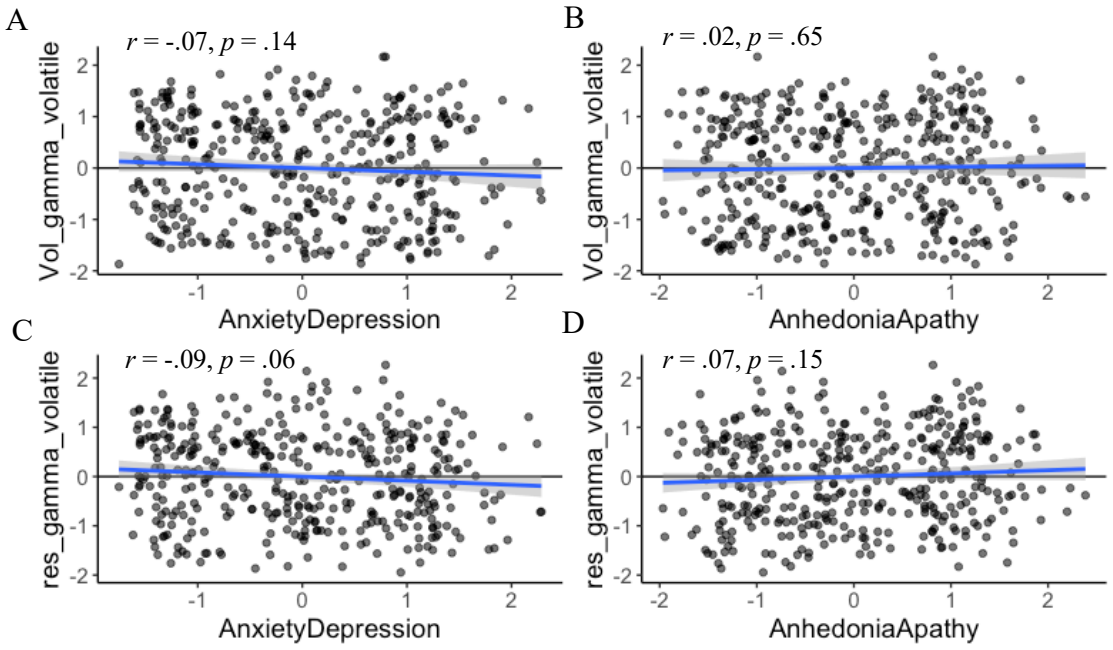

Figure 7.6. Partial correlation between factor scores and the gamma parameter in the volatile condition of the Volatility Task. (A) Correlation between Anxiety-Depression scores and the gamma parameter in the volatile condition (*Gamma-volatile*). (B) Correlation between Anhedonia-Apathy scores and Gamma-volatile. (C) and (D) show the same correlations after regressing out demographic variables from Gamma-volatile.

Higher Anhedonia-Apathy scores predicted a higher inverse temperature in the stable block (0.14, [0.03,0.25]), i.e. lower choice randomness. Without controlling for any variables Anhedonia-Apathy score significantly correlated with Beta-stable (r = .11, p = .02). After controlling for demographic variables, the relationship stayed identical (r = .11, p = .02). To test which subscale contributed to predicting Beta-stable, regressions were run separately for each subscale loading on Anhedonia-Apathy as regressor. No single subscale significantly predicted Beta-stable, but all regression weights had the sign congruent to their factor loading (i.e. AMI subscale predicted higher, and TEPS subscales lower Beta-stable). TEPS anticipatory (-0.09, [-0.19, 0.004]) and AMI social (0.09, [-0.0005, 0.19]) were closest to significance.

7.3.2.3 Probabilistic Instrumental Learning Task

We hypothesised that higher scores of internalising symptomatology (no factor-specific prediction) would predict lower performance in win trials (hypothesis 3). Higher Anxiety-Depression scores predicted lower performance in win (-0.13, [-0.23, -0.02]) as well as loss trials (-0.12, [-0.23, -0.01])(Figure 7.7). Without controlling for other variables, Anxiety-Depression scores were significantly negatively correlated to the performance in win trials (r = -.11, p = .01)(Figure 7.8A), and showed a non-significant negative correlation to performance in loss trials (r = -0.07, p = .14)(Figure 7.8B). After regressing out demographic variables from the performance measures, Anxiety-Depression scores still showed a significant negative correlation to performance in win trials (r = -.10, p = .03)(Figure 7.8C), and a trend towards a negative correlation to performance in loss trials (r = -.08, p = .09)(Figure 7.8D).

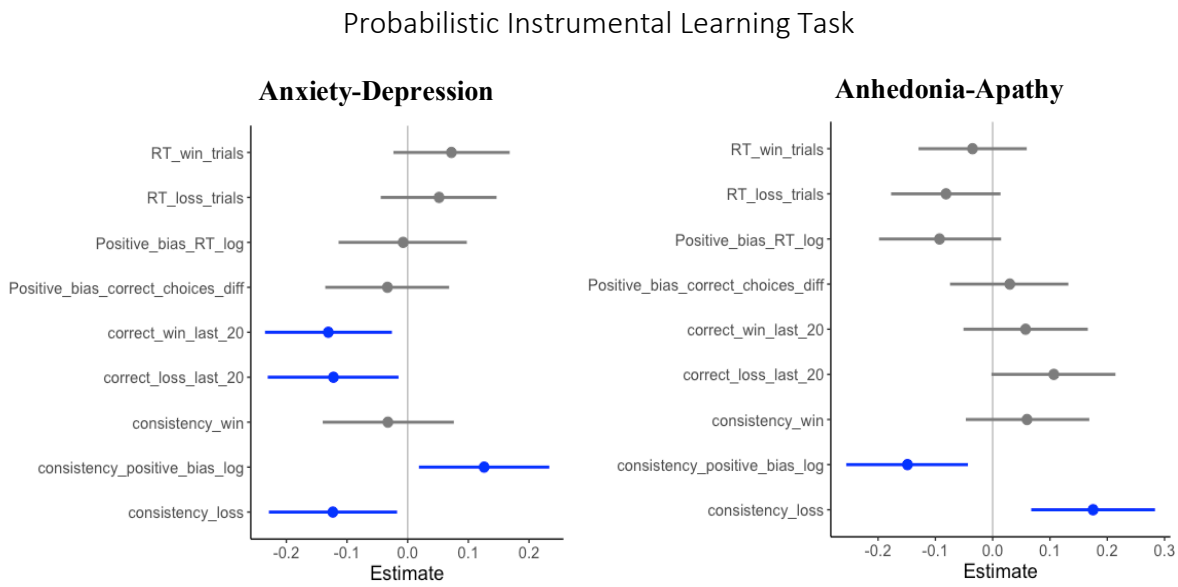

Figure 7.7. Regression weights for Anxiety-Depression and Anhedonia-Apathy for predicting individuals task measures extracted from the PILT. Error bars indicate 95% credible intervals.

To test which subscales contributed to predicting performance in win or loss trials, separate regressions with each subscale loading on Anxiety-Depression were run. The STICSA cognitive subscale significantly predicted lower performance in win trials (-0.13, [-0.24, -0.03]). No subscale alone predicted performance in loss trials. The regression weight for the STICSA somatic subscale was closest to significance (-0.10, [-0.20, 0.0007]).

Higher Anxiety-Depression scores predicted lower choice consistency for loss trials (-0.12, [-0.22, -0.01]) which led to an increase in positive consistency bias (i.e. higher consistency for win than for loss trials)(0.12, [0.01, 0.23])(Figure 7.7). In contrast, higher Anhedonia-Apathy predicted higher choice consistency in loss trials (0.17, [0.06, 0.28]) and lower positive consistency bias (-0.14, [-0.25,-0.04]). Without controlling for other variables, Anxiety-Depression scores showed a non-significant negative correlation with loss consistency (r = -.04, p = .38)(Figure 7.9A) and a non-significant positive correlation with the positive consistency bias (r = .03, p = .48). These correlations became somewhat stronger after controlling for demographic variables (consistency loss: r = -.06, p = .20, figure 7.9C; consistency positive bias log: r = .07, p = .14). Without controlling for other variables, Anhedonia-Apathy scores showed a significant positive correlation

with loss choice consistency (r = .15, p = .002)(Figure 7.9B) and a significant negative correlation with the positive consistency bias (r = -.12, p = .01). These correlations became slightly weaker but stayed significant after controlling for demographic variables (consistency loss: r = .12, p = .01, Figure 7.9D; consistency positive bias log: r = -.10, p = .04).

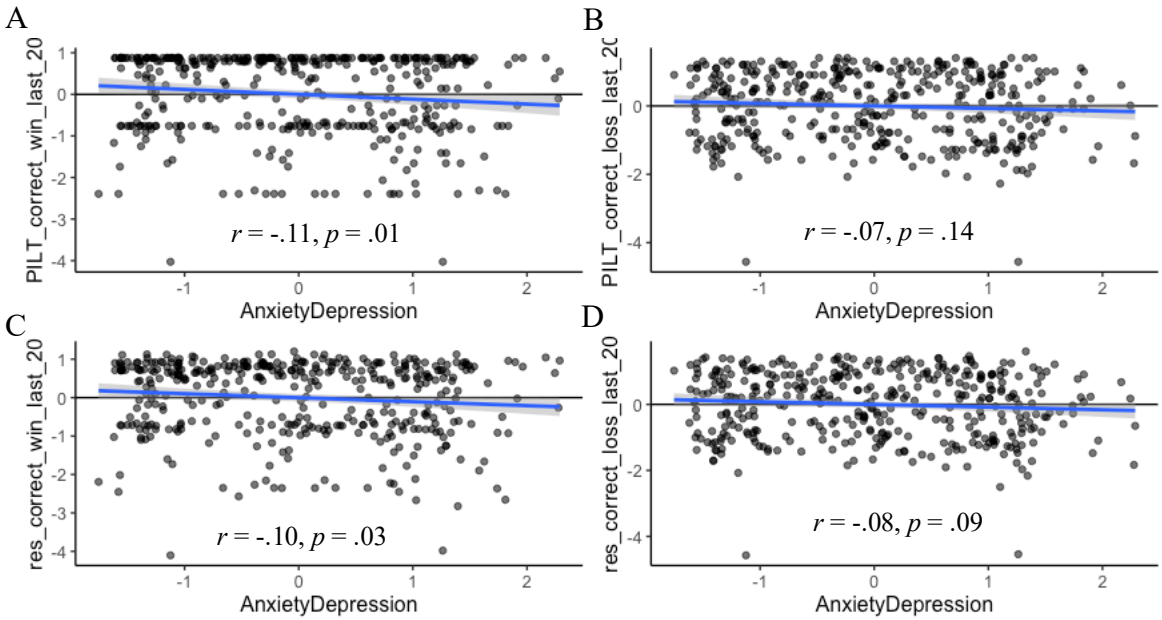

Figure 7.8. Partial correlations between Anxiety-Depression and performance in win trials (A,C) and performance in loss trials (B,D) in the PILT. (A) and (B) show the correlations between Anxiety Depression and performance in win or loss trials, respectively, without controlling for other variables. (C) and (D) show the same correlations after regressing out demographic variables from the performance measures.

No single subscale loading on Anxiety-Depression significantly predicted consistency in loss trials or positive consistency bias. The STICSA somatic subscale had the regression weight closest to significance (0.09, [-0.01, 0.19]) and predicted a higher positive bias in consistency. Regarding subscales loading on Anhedonia-Apathy, the AMI social subscale significantly predicted higher consistency in loss trials (0.11, [0.01, 0.20]). The TEPS anticipatory subscale significantly predicted lower consistency in loss trials (-0.11, [-0.21, -0.02]) and higher positive bias in consistency (0.10, [0.05, 0.002]) which is in line with the negative loading of the TEPS on the Anhedonia-Apathy factor.

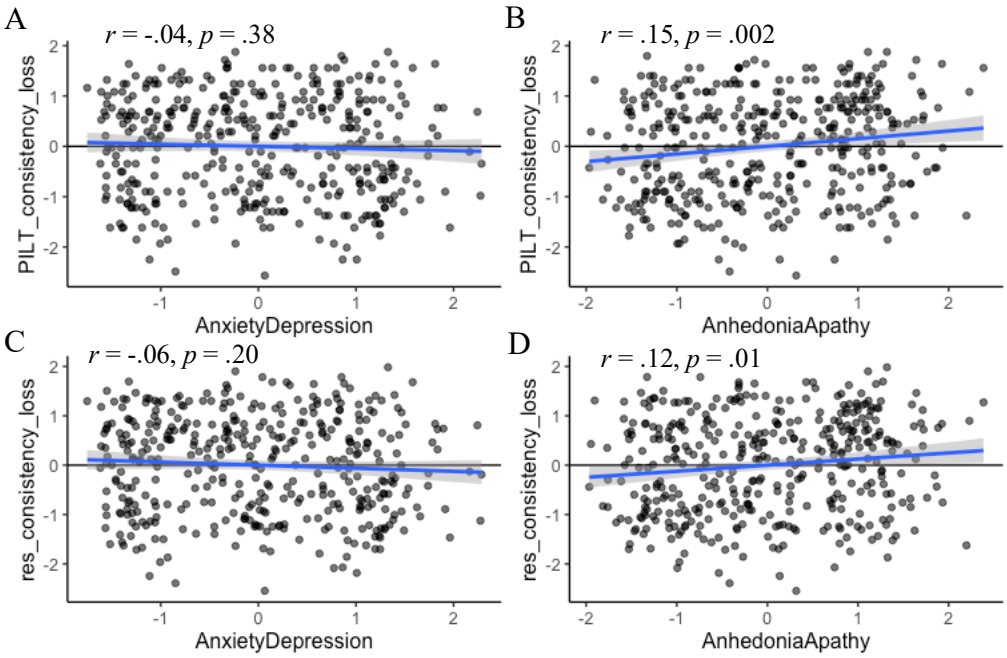

Figure 7.9. Correlation of Anxiety-Depression (A) and Anhedonia-Apathy scores (B) with choice consistency in loss trials in the PILT. (C) and (D) show the same correlations after regressing out demographic variables from the consistency measure.

7.3.2.4 Facial Expression Recognition Task

We hypothesised that stronger internalising symptomatology (no factor-specific prediction) would predict a lower positive bias in the unbiased hit rate and reaction times (hypothesis 4). In line with this, Anxiety-Depression scores significantly predicted lower positive bias in unbiased hit rate (-0.08, [-0.15, -0.01])(Figure 7.10A). Without controlling for other variables, Anxiety-Depression scores showed a non-significant negative correlation with the positive bias in unbiased hit rate (r = -.06, p = .16)(Figure 7.10B). After controlling for demographic variables, the correlation was slightly weaker (r = -.05, p = .28)(Figure 7.10C). Separate regressions for each subscale loading on Anxiety-Depression indicate that the relationship between the symptom factor scores and positive bias was mainly driven by the STICSA cognitive subscale (-0.12, [-0.19, -0.05]) and BDI (-0.07, [-0.14, -0.004]). Looking at individual emotions, the correlation between Anxiety-Depression and positive bias in unbiased hit rate seemed to be primarily driven by surprise (-0.07, [-0.14,0.01]) and disgust (0.07, [-0.02, 0.16])(although the regressors were not significant). Anxiety-Depression did

not significantly predict a positive bias in reaction times, but on a descriptive level the effect was in the same direction (-0.06, [-0.17, 0.05]).

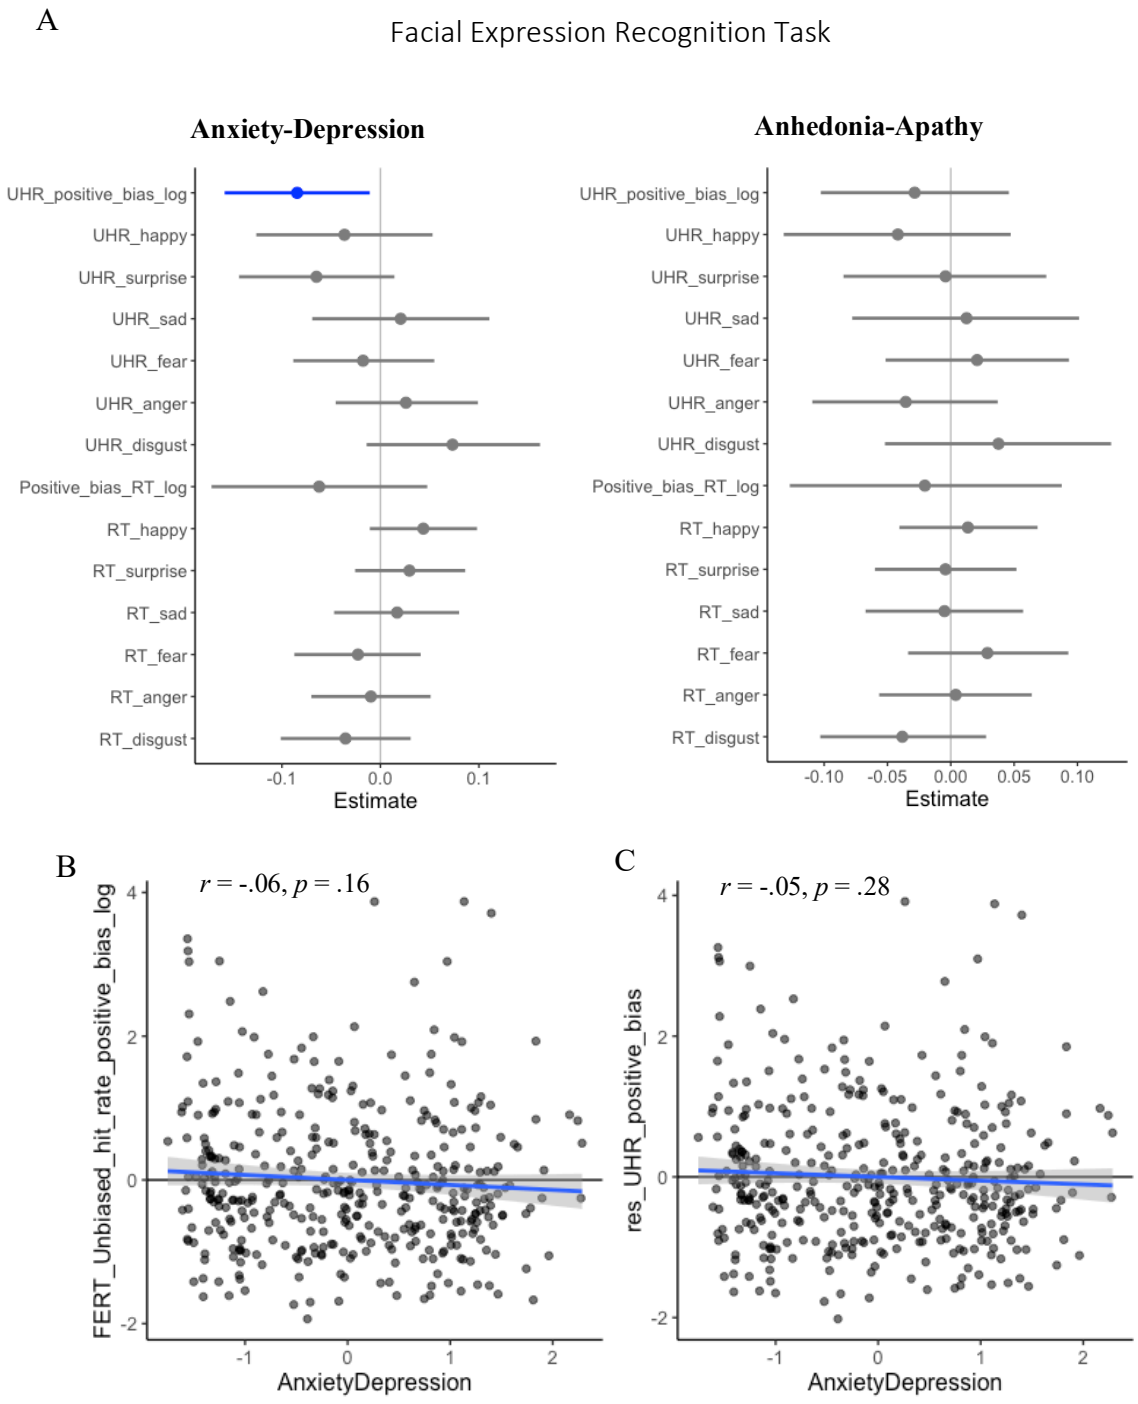

Figure 7.10. (A) Regression weights for Anxiety-Depression and Anhedonia-Apathy for predicting individuals task measures extracted from the FERT. Error bars indicate 95% credible intervals. (B) Partial correlations between positive bias in unbiased hit rate and Anxiety-Depression scores. (C) show the same correlation after regressing out demographic variables from the positive bias in unbiased hit rate.

7.4 Discussion

7.4.1 Summary of main findings

The aim of this study was to test whether parameters derived from five different tasks differentially correlate with distinct symptom dimensions of depression. In our exploratory dataset, we tested four main hypotheses. We hypothesised that apathy would predict stronger decision inertia in the Wheel of Fortune Task. However, there was no evidence for this relationship in our dataset. We further hypothesised that anxiety would predict lower learning rate adjustment in the Volatility Task. This hypothesis was not confirmed in our dataset. However, we found that Anxiety-Depression predicted less reliance on volatile loss probabilities. In line with our third hypothesis, symptoms related to anxiety and depression predicted lower performance in reward learning in the PILT. In line with our fourth hypothesis, symptom of anxiety and depression predicted a lower positive bias in accuracy in the FERT. In addition, some interesting findings emerged from exploratory analyses. In the PILT, Anxiety-Depression did not only predict lower performance in reward learning, but also in punishment learning which was characterised by lower choice consistency. In contrast, Anhedonia-Apathy predicted better performance and higher choice consistency in punishment learning. While Anxiety-Depression predicted lower weight on volatile loss probabilities in the Volatility Task, Anhedonia-Apathy predicted more weight on volatile loss probabilities and higher inverse temperature in the stable condition.

7.4.2 Factor analysis revealed two symptom dimensions

To reduce the dimensionality of the mood questionnaire items, we started by conducting exploratory factor analysis on the single items of STICSA, TEPS, AMI and BDI. This analysis largely recovered the questionnaire subscales reported in the literature. Previous studies have found varying numbers of factors for the BDI (Huang & Chen, 2015). In our study, all BDI items loaded on one factor. To further reduce the dimensionality, we conducted a second exploratory factor

analysis with oblique rotation on the questionnaire subscales (including one total score for the BDI) which resulted in two factors which were positively correlated. We labelled the first factor "Anxiety-Depression" since the two STICSA subscales and the BDI loaded on this factor. The second factor was labelled "Anhedonia-Apathy" since all subscales of the TEPS and AMI loaded on this factor. The AMI subscale "emotional sensitivity" loaded positively on the Anhedonia-Apathy factor, and negatively on the Anxiety-Depression factor. This suggests an interesting distinction between the two positively correlated factors: Individuals scoring high on Anxiety-Depression worry a lot about themselves and others, whereas individuals scoring high on Anhedonia-Apathy show a blunted emotional response. These two factors are roughly in line with the distinction of anxiety vs. anhedonia which has been suggested to differentiate different phenotypes of depression (Downar et al., 2014; Drysdale et al., 2017; Siddiqi et al., 2020b). Since we have only conducted exploratory factor analysis in one dataset, it is unclear how stable this two-factor solution is. However, previous research also reported factors related to anxiety vs. anhedonia (Clark, Steer, & Beck, 1994; Gagne et al., 2020; Steer, Clark, Beck, & Ranieri, 1995) which suggests that these factors are likely to reproduce in our validation dataset.

7.4.3 Discussion of expected and actual findings

The primary goal of this study was to explore relationships between these two symptom factors and cognitive task measures. Four main hypotheses were tested. First, we hypothesised that higher scores of apathy would predict higher decision inertia, i.e. a higher influence of the number of previous searches (Scholl et al., 2022). However, no such relationship was found, neither for Anhedonia-Apathy scores, nor for any of the underlying questionnaire subscales. In contrast, on a trend-level, individuals with higher Anhedonia-Apathy were more likely to stop searching as the prospective value of searching decreased, which contradicts the notion of decision inertia. Unexpectedly, we found that higher Anxiety-Depression predicted higher cost insensitivity, i.e. participants with higher Anxiety-Depression scores were less likely to be deterred by higher costs of spinning the wheel. Since previous research found this parameter to be correlated with

compulsivity (Scholl et al., 2022), we included the OCI-score in control analyses. Our results suggest that the correlation between Anxiety-Depression and cost insensitivity is driven by variance shared between Anxiety-Depression and OCI scores.

Second, we hypothesised that individuals with stronger anxiety symptoms would adjust their learning rate to a smaller extent to changes in volatility of punishment associations (Browning et al., 2015; Gagne et al., 2020). In contrast with this, Anxiety-Depression scores did not predict learning rate adjustment, neither did any of the underlying subscales. However, higher Anxiety-Depression scores predicted a lower weight on loss probabilities in the volatile condition. Although unexpected, this finding is compatible with the hypothesis that anxiety might be associated with difficulties adjusting learning rates to volatility. If highly anxious individuals struggle to track loss probabilities in the volatile condition, they might base their decisions primarily on loss magnitude instead. In a reinforcement learning model, this would be captured by a lower gamma parameter as observed in our study. Although our finding seems to be compatible with the results from previous studies, it is unclear why anxiety was related to a different model parameter although our paradigm and model were very similar to previous studies. Unexpectedly, we found that higher scores of Anhedonia-Apathy predicted a higher weight on loss probabilities in the volatile condition, and a higher inverse temperature in the stable condition. This suggests that individuals with higher Anhedonia-Apathy scores might be better at tracking loss probabilities and better at estimating the negative values of the two options than individuals with lower symptom levels. This is in line with previous research reporting that individuals with depression perform punishment learning tasks better than healthy individuals (Beevers et al., 2013). However, these findings were unexpected, and it remains to be seen if they replicate in the validation sample.

Third, we hypothesised that higher internalising symptomatology would predict lower performance in reward learning (Kumar et al., 2018; Walsh, Browning, et al., 2018). In line with this, higher scores of Anxiety-Depression predicted lower performance in reward learning in the PILT. Anxiety-Depression scores also predicted lower performance in punishment learning, indicating

that the relationship was not specific to the valence of the outcomes. However, Anxiety-Depression scores had a differential effect on choice consistency in reward vs. punishment learning. Higher Anxiety-Depression scores predicted lower choice consistency for punishment learning only which led to a higher positive bias in choice consistency (i.e. more consistent choices for reward than for punishment learning). This suggests that the lower performance in reward and punishment learning might be caused by different mechanisms. Low performance in punishment learning might be caused by a high number of switches between the options, whereas low performance in reward learning might be caused by choosing the incorrect shape more consistently. In computational terms, this might correspond to an increased punishment learning rate, and decreased reward learning rate, which are hypothesised to be a potential mechanism leading to negative biases in depression (Aylward et al., 2019; Pike & Robinson, 2022). In contrast, higher scores of Anhedonia-Apathy predicted higher choice consistency in punishment learning in the PILT. In line with this, Anhedonia-Apathy also predicted a lower positive bias in choice consistency, as well as a trend towards better performance in punishment learning. This is line with our findings from the Volatility task and with previous findings that individuals with depression might perform better in punishment learning tasks than healthy individuals (Beevers et al., 2013).

Fourth, we hypothesised that stronger internalising symptomatology would predict a lower positive bias in emotion recognition (Gur et al., 1992; Walsh, Browning, et al., 2018; Walsh, Huneke, et al., 2018). In line with this, higher Anxiety-Depression scores predicted a lower positive bias in choice accuracy in the FERT (as measured by unbiased hit rates). On a descriptive level, this seemed to be primarily driven by higher accuracy for surprise, and lower accuracy for disgust. Although Anxiety-Depression did not significantly predict lower positive bias in reaction times, on a descriptive level, the effect was in the same direction. Future analyses could combine accuracy and reaction times into one "efficiency" measure. Anhedonia-Apathy scores did not predict any of the measures derived from the FERT. Therefore, our results suggest that a reduced positive bias in emotion recognition is specific to Anxiety-Depression, which was mainly driven by the STICSA cognitive subscale and BDI score, i.e. cognitive symptoms.

7.4.4 Differentiating Anxiety-Depression from Anhedonia-Apathy

The critical feature of a system for defining different phenotypes of depression is its ability to differentiate between different phenotypes based on the measurements the system is based on (i.e. cognitive task markers in our case). Comparing the profiles of correlated task measures between Anxiety-Depression and Anhedonia-Apathy, two measures correlated with these symptom factors with an opposite sign and might therefore be particularly important for distinguishing between them. The first measure is Gamma-volatile in the Volatility Task which is negatively correlated with Anxiety-Depression, and positively with Anhedonia-Apathy. The second measure is choice consistency in punishment learning in the PILT, which is negatively correlated with Anxiety-Depression, and positively with Anhedonia-Apathy. Together, these two measures suggest that individuals scoring high on Anxiety-Depression might have difficulties tracking punishment probabilities, which was specific to volatile associations in the Volatility Task, but also present for stable associations in the PILT. Individuals scoring high on Anhedonia-Apathy, show higher reliance on probabilities vs. rewards in volatile environments, and are more consistent in punishment learning in stable environments, i.e. seem to perform better at punishment learning in general. Based on the exploratory and very preliminary analysis of this dataset, it therefore seems that the most informative cognitive process to distinguish between Anxiety-Depression and Anhedonia-Apathy might be punishment learning, with differences observed for volatile as well as stable associations.

7.4.5 Further steps

The primary aim of this project was to explore relationships between clinical symptoms and cognitive task measures. Across all analyses conducted so far, correlations between symptoms and task measures have been relatively small ($\leq .15$). The tasks included in this study have been selected based on the hypothesis that they might capture distinct cognitive processes altered in depression. As a next step, it would therefore be interesting to test how much variance in symptom

factors can be explained by all (or a selection of) task parameters together (i.e. use several task measures to predict scores of Anxiety-Depression or Anhedonia-Apathy). Comparing how much variance different combinations of task measures can explain might help to assess which task parameters capture distinct processes.

At later stages of the project, we are planning to apply more advanced analyses that aim at maximising the correlation between symptoms and task measures. More specifically, we are planning to use canonical correlation analysis (CCA) which aims at extracting linear combination of variables of two different modalities which maximally correlate (Hotelling, 1936; Winkler, Renaud, Smith, & Nichols, 2020). With respect to our study, this would correspond to linear combinations of symptoms that are maximally correlated with linear combinations of task measures. These "canonical variates" can be extracted in a way to be either orthogonal or, similar to factor analysis, to some degree correlated (Wang et al., 2020). In previous research, CCA has been used to relate clinical symptoms to resting-state functional connectivity (Drysdale et al., 2017; Smith et al., 2015), based on which potential phenotypes have been investigated (Drysdale et al., 2017). Since CCA is prone to overfitting, strict pipelines of permutation testing and out-of-sample validation need to be applied (Dinga et al., 2020; Dinga et al., 2019; Winkler et al., 2020). Suitable analysis pipelines are described in (Mihalik, Adams, & Huys, 2020; Wang et al., 2020).

7.4.6 Limitations

This study has several limitations. First, the sample is a non-clinical population recruited on an online platform. Although a substantial part of the participants indicated a current or past psychiatric diagnosis, it is unclear whether the findings would generalise to clinical populations. Second, in the exploratory analyses reported in this chapter, no correction for multiple comparisons was applied and most of the observed findings were unexpected. Therefore, it is likely that some findings might not replicate in an independent dataset. The exploratory dataset will be used for power analysis to ensure that the confirmation dataset will have adequate statistical power to detect

these effects. Third, there are many alternative tasks that could have been selected for the study. Further research is needed to test which combination of tasks can explain the highest proportion of variance in depressive symptoms. Fourth, this study does not allow for any conclusions about the clinical usefulness of the findings. The tasks will need to be included in clinical trials to test whether they have any predictive potential. Fifth, it needs to be tested whether the tasks are practically usable in the clinic. Although the fact that these tasks can be performed online increases accessibility, this requires the compliance of patients who might not be familiar with the performance of cognitive tasks. This also raises the question of whether there are any alternative measures that provide similar information but might be easier to use.

7.5 Conclusions

To conclude, we found that two symptom factors related to symptoms of anxiety/depression or anhedonia/apathy differentially correlated with specific parameters of four cognitive tasks in an exploratory dataset. The parameters that seemed to best differentiate between the two symptom factors were related to punishment learning in volatile as well as in stable environments. At later stages of the project, multivariate analyses will be applied aiming at maximising correlations between clinical symptoms and cognitive task markers. An independent validation dataset will be collected to test the hypotheses generated from the exploratory dataset.

7.6 Supplementary analyses

7.6.1 Exploratory factor analysis on single items of BDI, STICSA, TEPS and AMI

All 78 items from the BDI, STICSA, TEPS and AMI were entered into one exploratory factor analysis (using oblique rotation). Parallel analysis suggested an optimal solution of eight factors. The following plots show the loadings for each item. Bars with black filling are factor loadings above 0.4 which are considered significant. Red lines indicate items that should load (positively)

on the factor. The analysis largely recovered the subscales for each questionnaire (apart from BDI). All BDI items loaded on one factor.

BDI

All items loaded on one factor.

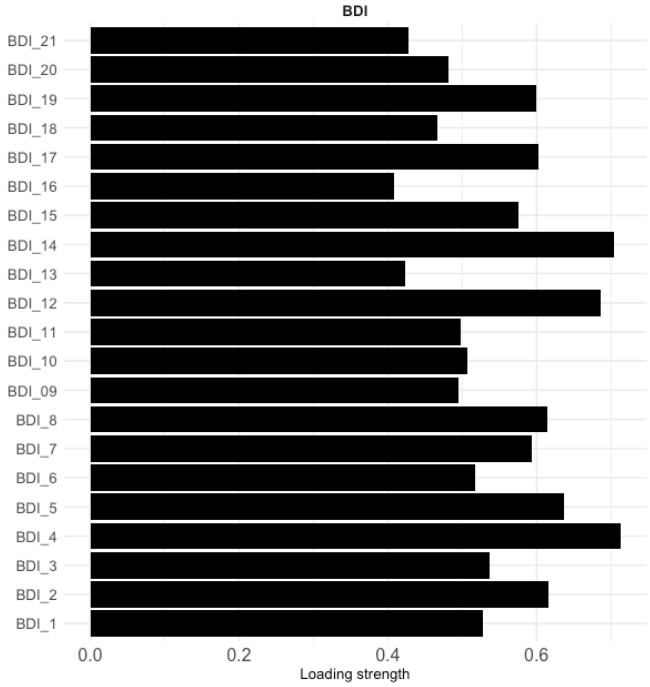

Figure 7.11. Loadings of the BDI items on the BDI factor.

STICSA

Two factors emerged for the STICSA. One factor was identical to the somatic anxiety factor reported in the literature. Items that should load on the cognitive anxiety factor also loaded on this factor in our analysis although most loadings were below 0.40.

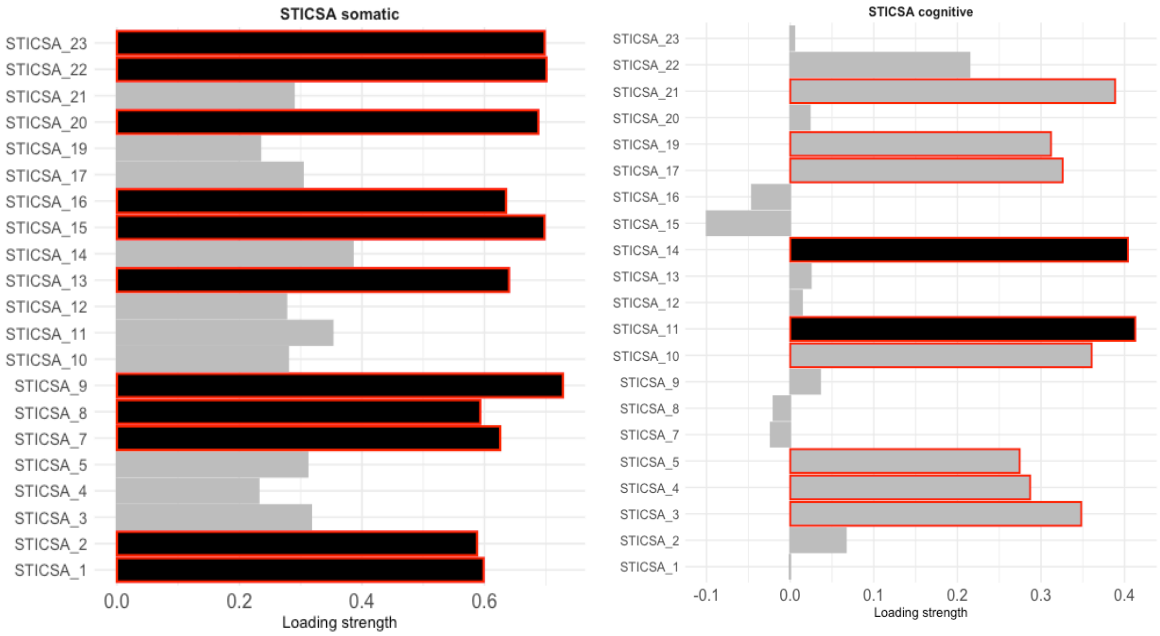

Figure 7.12. Loadings of the STICSA items on the somatic and cognitive factors.

TEPS

For the TEPS, two factors were extracted. These were largely in line with anticipatory and consummatory factors reported in the literature. Two items loaded negatively on the factor they should load on positively (but > -0.4); item 13 is: *"I really enjoy the feeling of a good yawn";* item 14 is: *"I don't look forward to things like eating out in restaurants"*, which was the only negatively phrased item which participants might have overlooked.

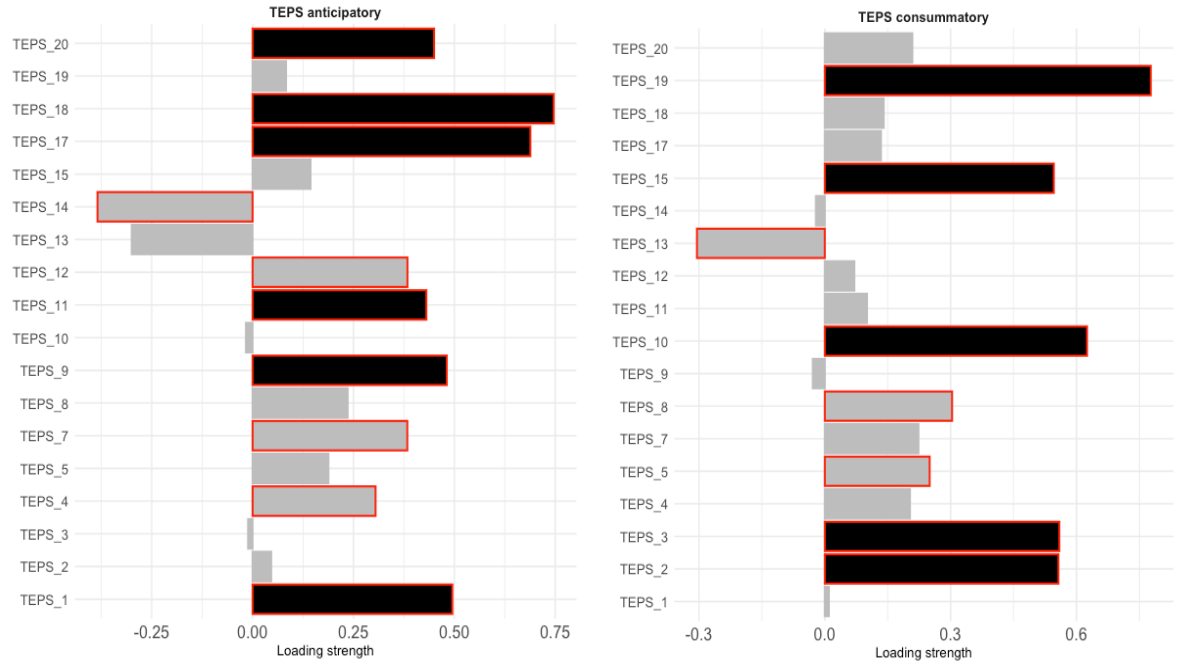

Figure 7.13. Loadings of the TEPS items on the anticipatory and consummatory factors.

AMI

Three factors were extracted which were very similar to the behavioural, social and emotional factors reported in the literature.

Figure 7.14. Loadings of the AMI items on the behavioural, social and emotional factors.

8 General discussion

In this thesis, two distinct projects are presented which investigated cognition-related approaches to the long-term goal of ultimately improving the use of NIBS in depression treatment. The first project used cognition as manipulation, while the second project used cognition as measure of depression psychopathology. The first project, comprising of three studies, set out to test whether tDCS applied during cognitive task performance might have the potential to reduce negative biases in depression. The second project was motivated by the question whether different phenotypes of depression could be identified based on cognitive task performance. This thesis includes preliminary analyses of an exploratory dataset testing which cognitive task parameters can capture distinct dimensions of depressive symptomatology.

8.1 Combining tDCS and reinforcement learning to normalise learning rate adjustment

The initial rationale of this proof-of-concept project in healthy individuals was to investigate whether tDCS might be able to increase the relative balance of reward vs. punishment learning, with the ultimate translational goal that such an approach could potentially be used to counter-act negative biases in depression. In healthy individuals, tDCS did indeed increase reward learning specifically (chapter 4). However, this effect was only present in conditions where wins were uninformative. To test whether such an effect might be useful to normalise information processing in depression, we compared choice behaviour in this same task between individuals with low mood and healthy controls (chapter 5). In contrast to our predictions, individuals with low mood did not show increased punishment vs. reward learning. However, they did differ in one specific parameter which captured the extent to which individuals adjust their win vs. loss learning rates to changes in informativeness. More specifically, individuals with low mood adjusted their loss learning rate less, and their win learning rate more than healthy individuals to changes in informativeness. The

third study set out to test whether the effect of bifrontal tDCS observed in chapter 4 (i.e. increase in reward learning) would replicate in individuals with low mood (chapter 6). In contrast to our hypothesis, tDCS did not increase reward learning in low mood. However, there was evidence that tDCS might normalise learning rate adjustment in low mood.

8.1.1 Reconsidering reinforcement learning in depression

The rationale of this project was based on the hypothesis that negative biases in depression might arise from increased punishment vs. reward learning (Aylward et al., 2019; Eshel & Roiser, 2010; Pike & Robinson, 2022). The initial aim of this project was therefore to increase reward learning (or decrease punishment learning) which might be useful to counter-act negative biases in depression. In healthy participants, tDCS indeed increased reward learning, which we thought might potentially be helpful in the treatment of depression. Following the same logic, our lab also previously conducted a training study in healthy volunteers, in which the same experimental paradigm was transformed into a positive bias training task, by repeatedly presenting the condition in which wins were more informative (Overman, Browning, & O'Shea, 2020). Participants showed increased reward learning rates with training, but this effect did not change performance on a transfer task of emotional face processing. Bifrontal tDCS had no effect on behaviour in the training task or on the transfer task. While the idea of increasing reward vs. punishment learning shaped the study design for this thesis, our findings suggest a different perspective on reinforcement learning in depression.

Comparing task performance between individuals with low mood and healthy controls, we found that low mood was characterised by alterations in learning rate adjustment. Individuals with low mood adjusted their loss learning rate less, and their win learning rate more to changes in informativeness than healthy individuals. Reduced loss learning rate adjustment is in line with previous studies reporting difficulties in learning rate adjustment associated with depression (Browning et al., 2015; Gagne et al., 2020). The increase in win learning rate adjustment was

unexpected, and might be a compensatory strategy. Difficulties in adjusting learning rates appropriately to changes in informativeness has been suggested as a potential information processing deficit leading to negative biases (Pulcu & Browning, 2017). However, in our study there was no evidence for a negative bias *per se* (e.g. increased punishment vs. reward learning rates).

We concluded from this study (chapter 5) that (a tDCS-induced) increase in reward learning is unlikely to normalise reinforcement learning deficits in depression. Rather than teaching individuals to focus on positive information, our data suggest that we should instead train them to adjust their learning rates appropriately to the environment. We chose the IBLT as experimental paradigm to test how tDCS influences learning rates for wins and losses depending on their relative informativeness. At the same time, the IBLT could be used as a training paradigm. By performing the IBLT repeatedly, individuals would be trained in adjusting their learning rates appropriately to changes in informativeness, i.e. they would be trained in a cognitive ability which appears to be impaired in depression. In our third study, we found that tDCS applied during the IBLT normalised learning rate adjustment in low mood. Although this finding was unexpected, this suggests that tDCS applied during this task might have a potentially beneficial effect on reinforcement learning in depression.

8.1.2 Reconsidering the effect of tDCS

The bifrontal tDCS setup was primarily chosen to increase excitability in the left DLPFC, an area which is hypoactive in depression (George et al., 1994). We hypothesised that bifrontal tDCS which increases excitability in the left DLPFC and decreases excitability in the right DLPFC would increase the processing of positive vs. negative information. The increase in reward learning observed in our first study supports this hypothesis. However, in individuals with low mood, bifrontal tDCS did not increase reward learning. Instead, bifrontal tDCS normalised learning rate adjustment in individuals with low mood. It is unclear why the effect of bifrontal tDCS differed

between healthy individuals and individuals with low mood. The effects of tDCS have been shown to depend on baseline behaviour (Clarke et al., 2014). Comparing task performance between healthy individuals and individuals with low mood during sham tDCS, we found that individuals with low mood showed higher switch probabilities, increased choice randomness, and altered learning rate adjustment (chapter 5). The most plausible explanation may therefore be that individuals with low mood perform the task in a different way than healthy participants which might lead to a different effect of tDCS.

As discussed above, the IBLT could be used as a training paradigm to train participants in adjusting their learning rates to different volatility conditions. TDCS has been shown to enhance synaptic plasticity and learning effects (Fritsch et al., 2010; Reis et al., 2009). Applying tDCS during repeated performance of the IBLT might therefore increase training effects and help participants to improve in adjusting their learning rates appropriately.

8.1.3 Clinical relevance

The most important question arising from this project is whether the results are potentially clinically meaningful or useful. The finding that tDCS might normalise learning rate adjustment in low mood (chapter 6) seems promising, however, it remains to be tested whether this would lead to symptom improvement. Whether the behavioural effect observed in this study might lead to clinical benefits depends on two aspects, *generalisability* and *causality*. It is unclear whether this effect observed in our study would generalise to other contexts (tasks) and ultimately to real-world scenarios. It is difficult to imagine that training effects in a task as abstract as the IBLT would transfer to complex decision-making in real life. However, cognitive bias modification, which uses abstract training tasks to counteract negative biases, has been found to reduce symptoms of depression and anxiety (although effect sizes tend to be small)(Fodor et al., 2020). More recent research is aiming at developing "ecologically valid" paradigms designed to capture cognitive processes particularly relevant in real life (Scholl & Klein-Flugge, 2018). But even if the effects

observed in this study generalise to real life, it remains to be tested whether they would actually lead to symptom improvement. The rationale of targeting negative biases in depression is based on the hypothesis that negative biases *cause* depressive symptoms. Some evidence indeed suggests that reducing negative biases might lead to symptom improvement (Browning et al., 2019; Fodor et al., 2020; Godlewska et al., 2016). In our study, tDCS had no effect on any negative bias measure per se. However, it had an influence on learning rate adjustment which is hypothesised to be a potential mechanism leading to negative biases (Pulcu & Browning, 2017, 2019). Future studies need to test for a causal relationship between learning rate adjustment and clinical symptoms. This could be tested in a future training study in which individuals suffering from depression receive daily training in learning rate adjustment (e.g. using the IBLT). It should be tested whether the training effects transfer to other tasks, and whether improvements in the ability to adjust learning rates eventually reduce depressive symptoms.

8.1.4 Combining tDCS with neuroimaging

While the effect of bifrontal tDCS in individuals with low mood had an effect that might potentially be relevant in depression treatment (chapter 6), the study does not allow for conclusions regarding the physiological basis of this effect. The bifrontal tDCS setup stimulates large areas in the brain, and we can only speculate which brain areas the behavioural effect relies on. The DLPFC and the medial prefrontal cortex (which might receive the highest electric field strength (Karabanov et al., 2019)) are part of a network associated with reinforcement learning processes (Farashahi et al., 2019; Massi et al., 2018; Rushworth, 2008). One key region activated in response to volatility is the dorsal anterior cingulate cortex (Behrens et al., 2007). The effect of tDCS on learning rate adjustment is therefore likely to rely on functional changes in prefrontal cortical regions as well as the dorsal anterior cingulate cortex. Understanding the neural mechanism of the behavioural effect is important to optimise the electrode setup and maximise the behavioural effect. A better understanding of the physiological basis could be gained by combining tDCS with fMRI. As a first step, structural scans could be used to calculate each individual's electric field

distribution across the brain. To test which brain region the behavioural effect relies on, the strength of the behavioural effect could be correlated to the electric field strength in different brain regions across individuals. This approach would be relatively simple to implement, but relies on electric field modelling. To directly test for tDCS effects on functional brain activity, fMRI scans could be recorded during stimulation while participants perform the task. The functional activation (and connectivity) of different brain regions could be compared between real and sham tDCS to test for effects on functional activation in different brain regions. Moreover, the strength of the behavioural effect could be correlated to the strength of the effect on functional activation in different brain regions across participants. Insights about the physiological basis of the behavioural effect could then be used to further optimise the tDCS setup, for example by using a more focal multi-electrode array targeting a specific brain region (Poydasheva et al., 2021).

8.1.5 Applying tDCS during a learning task

One important consideration that this project was built on is the rationale that tDCS induces plasticity and learning (Fritsch et al., 2010) and might therefore be more effective if applied during a learning task. The effects of tDCS on choice behaviour were only observed for online tDCS, i.e. when the stimulation was applied during task performance. This supports the idea that tDCS enhances activity-dependent synaptic plasticity. Future clinical trials should test whether applying tDCS during a cognitive task might have stronger antidepressant effects than applying tDCS at rest. The key question is which learning task would be ideal for this purpose. For this project, the IBLT was chosen because of the recent findings suggesting that alterations in reinforcement learning might lead to negative biases in depression. However, negative biases might arise at different levels in the information processing hierarchy, ranging from low-level attentional processes (Bishop, Duncan, Brett, & Lawrence, 2004), through information integration (Gagne et al., 2022; Sharot et al., 2011) to interpretation (Daniel-Watanabe et al., 2020). It might be necessary to target negative biases on multiple hierarchical levels to lead to clinically meaningful changes. This might require the combination of different interventions, for instance brain

stimulation to target lower levels, and psychotherapy to target higher levels of information processing.

8.2 Cognitive phenotyping – first insights and challenges

The field of NIBS depression treatment is moving towards personalising the target region. As one approach of personalising the target region phenotyping has been suggested, i.e. finding each individual's depression phenotype and stimulate the region this phenotype best responds to. Previous research indicates that patients who respond to different target regions differ in the relative severity of anxiety vs. anhedonia symptoms (Downar et al., 2014; Drysdale et al., 2017). Since phenotyping purely based on clinical symptoms has not been found to predict treatment response reliably (Arnow et al., 2015; Uher et al., 2011), and neuroimaging would be difficult to implement in clinical practice, we are interested in the question whether phenotypes of depression could be identified based on cognitive task performance. Since a variety of cognitive tasks have been shown to relate to depressive symptoms, the aim of this study was to inform the selection of tasks which might be included in clinical trials in future stages of the project. Since the relative severity of anxiety and anhedonia seems to have predictive potential, the study was set out to test which parameters of five cognitive tasks (four of which have been analysed so far) differentially relate to symptoms of anxiety and anhedonia.

8.2.1 Relating anxiety and anhedonia to cognitive task parameters

To reduce the dimensionality of the mood questionnaires, we started by conducting exploratory factor analysis on the questionnaire subscales. This analysis yielded two factors which we labelled "Anxiety-Depression" (STICSA subscales and BDI) and "Anhedonia-Apathy" (TEPS and AMI subscales). In the preliminary analyses presented in this thesis, we tested whether the key findings from previous studies replicated in our dataset. Moreover, we explored relationships between the symptom factors and other relevant task parameters.

For two of the tasks, we did not observe the previously reported relationship between task parameters and psychiatric symptoms. In the Wheel of Fortune Task, symptoms of apathy have previously been shown to predict higher decision inertia, i.e. searching the alternatives for longer than appropriate given the value of the alternatives (Scholl et al., 2022). We did not observe this relationship in our study. If anything, the results pointed towards a relationship in the opposite direction. For the Volatility Task, it has previously been found that anxiety is related to lower adjustment of learning rates to changes in informativeness (Browning et al., 2015; Gagne et al., 2020). No such relationship was present in our dataset. However, we found that Anxiety-Depression predicted a lower weight on loss probabilities when loss associations were volatile. One possible explanation for this might be that highly anxious individuals have difficulties tracking the volatile loss probabilities and as a compensatory strategy place more weight on loss magnitude. This would support the hypothesis that anxiety is associated with difficulties tracking volatile loss associations. Although this explanation sounds plausible, it is unclear why anxiety was related to lower learning rate adjustment in previous studies using the identical (or very similar) paradigm, and with lower weight on probabilities in this study.

The main hypotheses for the other two tasks were confirmed. Anxiety-Depression predicted a negative bias in the FERT, and lower performance in the PILT. Moreover, some unexpected relationships were observed in the dataset. In contrast to Anxiety-Depression, Anhedonia-Apathy predicted higher inverse temperature in the stable condition and higher weight on loss probabilities in the volatile condition in the Volatility Task. While Anxiety-Depression predicted lower performance in reward and punishment learning in the PILT, Anhedonia-Apathy predicted higher performance in the punishment learning condition. Taken together, the results suggest that punishment learning in volatile as well as stable environments is the cognitive process which can best distinguish between the two symptom factors. The simplest explanation for this pattern of results is that individuals with high Anxiety-Depression scores generally perform worse in punishment learning, whereas individuals with high Anhedonia-Apathy generally perform better in punishment learning.

The most important (and potentially most difficult) step in developing a system for cognitive phenotyping is task selection. While this first study included a selection of five tasks, many alternative tasks capturing other relevant cognitive processes could be included. For instance, depressive symptoms have been linked to alterations in meta-cognition (Rouault et al., 2018), temporal discounting (Pulcu et al., 2014) or (cognitive) effort (Horne, Topp, & Quigley, 2021). Including additional tasks capturing other cognitive processes affected in depression might help to explain more variance in depressive symptoms.

8.2.2 Limitations and future directions

This study has several limitations. First of all, for this thesis, only the exploratory dataset has been collected. The purpose of the exploratory dataset is to explore the correlation structure between task measures and symptoms, and find the optimal methods for data analysis. Hypotheses generated from the exploratory dataset will be tested in a second independent dataset (yet to be collected) to test if the findings replicate.

Second, the relationships between psychiatric symptoms and task parameters observed in this study are very small. Previous studies found similarly low correlations between symptoms and task parameters (Cyders & Coskunpinar, 2011; Eisenberg et al., 2019; McHugh et al., 2011). Some authors even concluded that task parameters might not be meaningfully related to real-life outcomes (Eisenberg et al., 2019). This issue can be addressed in different ways. First, careful consideration needs to be given when defining the task parameters of interest. For example, many tasks include reaction time recordings, and reaction times between different tasks tend to be positively correlated (i.e. participants who respond quickly in one task are likely to also respond quickly in other tasks). However, raw reaction times might not be related to psychiatric symptoms in any meaningful way (they might relate to control variables such as age). A measure derived from reaction times that might be more relevant to psychiatry would for example be the ratio of reaction times towards negative vs. positive stimuli, i.e. a measure of positive bias in reaction

times. Careful inspection of the correlation matrix between measures derived from different tasks is necessary to investigate how different task parameters relate to one another and which of them might capture processes relevant to psychiatry. Second, analyses aiming at maximising correlations between task measures and symptoms, such as canonical correlation analysis, can be used. The correlations between linear combinations of task parameters and psychiatric symptoms extracted in canonical correlation analysis should be much higher than correlations between single task parameters and symptom factors reported in this thesis. However, this type of analysis requires a sophisticated analysis pipeline to assess significance and reliability. Third, one reason for small relationships between symptoms and task parameters might be low reliability of the task measures. A selection criterion for task measures should therefore be sufficient test-retest reliability which can unfortunately not be determined in our dataset since it only includes one testing timepoint. Fourth, alternative outcome measures should be considered. While self-reported symptoms are very commonly used as outcome measure, functional outcomes, such as someone's ability to work, might be equally (or even more) important. Cognitive task parameters might be more strongly related to functional outcome measures than to self-reported symptoms, and might be able to predict functional outcomes beyond self-reported symptom scores.

Third, it is unclear if the task parameters included in this study have any predictive potential for treatment response. This is especially the case given the weak relationships between task parameters and symptoms. To investigate whether the performance of cognitive tasks has predictive values, they need to be included in clinical trials. To test whether cognitive task measures can predict treatment response, the tasks could be performed prior to the intervention (e.g. TMS treatment). In addition, the tasks could also be included during and after the intervention period to investigate how symptoms and cognitive processes change in response to the treatment.

Fourth, symptoms and parameters are measured on a continuous scale, but need to be translated into discrete categories to guide treatment selection. To test for the existence of distinct phenotypes, clustering analysis could be applied to dimensions of symptoms and cognitive

parameters similar to (Drysdale et al., 2017). Future studies will then need to test whether different phenotypes preferentially respond to different treatments. However, it may very well be that clinical symptoms and task parameters are best described as continuous distributions without the existence of clusters. Alternatively, a classifier could be trained on task measures and clinical symptoms to predict treatment response.

8.3 Concluding remarks: Relevance of cognition in NIBS treatment

This thesis investigated two cognitive approaches to improving NIBS application in the treatment of depression. Although depression is characterised by many cognitive symptoms, surprisingly little research on NIBS depression treatment has focused on cognition. Over the past decade, useful insights into the neural mechanisms of NIBS treatment have been gained using neuroimaging. However, little research is available into cognitive mechanisms linking these neural mechanisms to clinical symptoms. In this thesis, interesting insights into cognitive mechanisms and potential clinical applications have been gained. We found that depression is associated with decreased adjustment of punishment learning rates and increased adjustment of reward learning rates, and tDCS might be able to normalise this deficit. As a novel treatment approach, we suggest to apply tDCS during the performance of a reinforcement learning task which might have larger antidepressant effects than tDCS alone. We have also begun to explore relationships between cognitive task parameters and distinct dimensions of depressive symptoms, in the hope that such work will in future yield insights that could prove useful in clinical trials for investigating the predictive value of cognitive task parameters for personalising NIBS treatment.

Contributions

We would like to thank Jacqueline Scholl and Hailey Trier for providing code, protocols and analysis scripts for the Wheel of Fortune Task and Foraging Fish Game. We would like to thank Amy Gillespie for providing task protocols and analysis scripts for the Facial Expression Recognition Task and Probabilistic Instrumental Learning Task (Gillespie, A., Gould van Praag, C. D., de Cates, A., Kaltenboeck, A., Lin, W., Howard, W., … Kretschmer, C. R. (2022, November 24). PERL Gorilla Tasks. Retrieved from osf.io/kj5r4).

The authors would like to acknowledge the use of the University of Oxford Advanced Research Computing (ARC) facility in carrying out this work. http://dx.doi.org/10.5281/zenodo.22558

Literature

Albein-Urios, N., Chase, H., Clark, L., Kirkovski, M., Davies, C., & Enticott, P. G. (2019). Increased perseverative errors following high-definition transcranial direct current stimulation over the ventrolateral cortex during probabilistic reversal learning. *Brain Stimul, 12*(4), 959-966. doi:10.1016/j.brs.2019.02.013

Alexander, W. H., & Brown, J. W. (2011). Medial prefrontal cortex as an action-outcome predictor. *Nat Neurosci, 14*(10), 1338-1344. doi:10.1038/nn.2921

Alonzo, A., Fong, J., Ball, N., Martin, D., Chand, N., & Loo, C. (2019). Pilot trial of home-administered transcranial direct current stimulation for the treatment of depression. *J Affect Disord, 252*, 475-483. doi:10.1016/j.jad.2019.04.041

American Psychiatric Association, D., & Association, A. P. (2013). *Diagnostic and statistical manual of mental disorders: DSM-5* (Vol. 5): American psychiatric association Washington, DC.

Ang, Y. S., Lockwood, P., Apps, M. A., Muhammed, K., & Husain, M. (2017). Distinct Subtypes of Apathy Revealed by the Apathy Motivation Index. *PLoS One, 12*(1), e0169938. doi:10.1371/journal.pone.0169938

Arfai, E., Theano, G., Montagu, J., & Robin, A. (1970). A controlled study of polarization in depression. *The British Journal of Psychiatry, 116*(533), 433-434.

Arnow, B. A., Blasey, C., Williams, L. M., Palmer, D. M., Rekshan, W., Schatzberg, A. F., . . . Rush, A. J. (2015). Depression Subtypes in Predicting Antidepressant Response: A Report From the iSPOT-D Trial. *Am J Psychiatry, 172*(8), 743-750. doi:10.1176/appi.ajp.2015.14020181

Aylward, J., Valton, V., Ahn, W. Y., Bond, R. L., Dayan, P., Roiser, J. P., & Robinson, O. J. (2019). Altered learning under uncertainty in unmedicated mood and anxiety disorders. *Nat Hum Behav, 3*(10), 1116-1123. doi:10.1038/s41562-019-0628-0

Baeken, C., Duprat, R., Wu, G. R., De Raedt, R., & van Heeringen, K. (2017). Subgenual Anterior Cingulate-Medial Orbitofrontal Functional Connectivity in Medication-Resistant Major Depression: A Neurobiological Marker for Accelerated Intermittent Theta Burst Stimulation Treatment? *Biol Psychiatry Cogn Neurosci Neuroimaging, 2*(7), 556-565. doi:10.1016/j.bpsc.2017.01.001

Baeken, C., Marinazzo, D., Wu, G. R., Van Schuerbeek, P., De Mey, J., Marchetti, I., . . . De Raedt, R. (2014). Accelerated HF-rTMS in treatment-resistant unipolar depression: Insights from subgenual anterior cingulate functional connectivity. *World J Biol Psychiatry, 15*(4), 286-297. doi:10.3109/15622975.2013.872295

Baller, E. B., Kaczkurkin, A. N., Sotiras, A., Adebimpe, A., Bassett, D. S., Calkins, M. E., . . . Satterthwaite, T. D. (2020). Neurocognitive and functional heterogeneity in depressed youth. *Neuropsychopharmacology*. doi:10.1038/s41386-020-00871-w

Barbieri, M., Negrini, M., Nitsche, M. A., & Rivolta, D. (2016). Anodal-tDCS over the human right occipital cortex enhances the perception and memory of both faces and objects. *Neuropsychologia, 81*, 238-244. doi:10.1016/j.neuropsychologia.2015.12.030

Beck, A. T. (1967). *Depression: Clinical, experimental, and theoretical aspects*: Hoeber Medical Division, Harper & Row.

Beck, A. T. (2002). Cognitive models of depression. *Clinical advances in cognitive psychotherapy: Theory and application, 14*(1), 29-61.

Beck, A. T., Steer, R. A., & Brown, G. K. (1996). Manual for the beck depression inventory-II. In: San Antonio, TX: Psychological Corporation.

Beevers, C. G., Worthy, D. A., Gorlick, M. A., Nix, B., Chotibut, T., & Todd Maddox, W. (2013). Influence of depression symptoms on history-independent reward and punishment processing. *Psychiatry Res, 207*(1-2), 53-60. doi:10.1016/j.psychres.2012.09.054

Behrens, T. E., Woolrich, M. W., Walton, M. E., & Rushworth, M. F. (2007). Learning the value of information in an uncertain world. *Nat Neurosci, 10*(9), 1214-1221. doi:10.1038/nn1954

Beijers, L., Wardenaar, K. J., van Loo, H. M., & Schoevers, R. A. (2019). Data-driven biological subtypes of depression: systematic review of biological approaches to depression subtyping. *Mol Psychiatry, 24*(6), 888-900. doi:10.1038/s41380-019-0385-5

Bergmann, T. O., Groppa, S., Seeger, M., Molle, M., Marshall, L., & Siebner, H. R. (2009). Acute changes in motor cortical excitability during slow oscillatory and constant anodal transcranial direct current stimulation. *J Neurophysiol, 102*(4), 2303-2311. doi:10.1152/jn.00437.2009

Berlim, M. T., Van den Eynde, F., & Jeff Daskalakis, Z. (2013). Clinically meaningful efficacy and acceptability of low-frequency repetitive transcranial magnetic stimulation (rTMS) for treating primary major depression: a meta-analysis of randomized, double-blind and sham-controlled trials. *Neuropsychopharmacology, 38*(4), 543-551. doi:10.1038/npp.2012.237

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. *Psychol Med, 44*(2), 225-239. doi:10.1017/S0033291713000512

Berman, M. G., Peltier, S., Nee, D. E., Kross, E., Deldin, P. J., & Jonides, J. (2011). Depression, rumination and the default network. *Soc Cogn Affect Neurosci, 6*(5), 548-555. doi:10.1093/scan/nsq080

Betancourt, M., & Girolami, M. (2015). Hamiltonian Monte Carlo for hierarchical models. *Current trends in Bayesian methodology with applications, 79*(30), 2-4.

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., 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. *J Physiol, 557*(Pt 1), 175-190. doi:10.1113/jphysiol.2003.055772

Bindman, L. J., Lippold, O., & Redfearn, J. (1962). Long-lasting changes in the level of the electrical activity of the cerebral cortex produced by polarizing currents. *Nature, 196*(4854), 584-585.

Bishop, S., Duncan, J., Brett, M., & Lawrence, A. D. (2004). Prefrontal cortical function and anxiety: controlling attention to threat-related stimuli. *Nat Neurosci, 7*(2), 184-188. doi:10.1038/nn1173

Bishop, S. J., & Gagne, C. (2018). Anxiety, depression, and decision making: a computational perspective. *Annual review of neuroscience, 41*, 371-388.

Blumberger, D. M., Vila-Rodriguez, F., Thorpe, K. E., Feffer, K., Noda, Y., Giacobbe, P., . . . Downar, J. (2018). Effectiveness of theta burst versus high-frequency repetitive transcranial magnetic stimulation in patients with depression (THREE-D): a randomised non-inferiority trial. *The Lancet, 391*(10131), 1683-1692. doi:10.1016/s0140-6736(18)30295-2

Bradley, B. P., Mogg, K., & Williams, R. (1995). Implicit and explicit memory for emotion-congruent information in clinical depression and anxiety. *Behaviour research and therapy, 33*(7), 755-770.

Brakowski, J., Spinelli, S., Dorig, N., Bosch, O. G., Manoliu, A., Holtforth, M. G., & Seifritz, E. (2017). Resting state brain network function in major depression - Depression symptomatology, antidepressant treatment effects, future research. *J Psychiatr Res, 92*, 147-159. doi:10.1016/j.jpsychires.2017.04.007

Brown, V. M., Zhu, L., Solway, A., Wang, J. M., McCurry, K. L., King-Casas, B., & Chiu, P. H. (2021). Reinforcement Learning Disruptions in Individuals With Depression and Sensitivity to Symptom Change Following Cognitive Behavioral Therapy. *JAMA Psychiatry, 78*(10), 1113-1122. doi:10.1001/jamapsychiatry.2021.1844

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. *Nat Neurosci, 18*(4), 590-596. doi:10.1038/nn.3961

Browning, M., Kingslake, J., Dourish, C. T., Goodwin, G. M., Harmer, C. J., & Dawson, G. R. (2019). Predicting treatment response to antidepressant medication using early changes in emotional processing. *Eur Neuropsychopharmacol, 29*(1), 66-75. doi:10.1016/j.euroneuro.2018.11.1102

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. *Br J Psychiatry, 208*(6), 522-531. doi:10.1192/bjp.bp.115.164715

Brunoni, A. R., Moffa, A. H., Sampaio-Junior, B., Borrione, L., Moreno, M. L., Fernandes, R. A., . . . Investigators, E.-T. (2017). Trial of Electrical Direct-Current Therapy versus Escitalopram for Depression. *N Engl J Med, 376*(26), 2523-2533. doi:10.1056/NEJMoa1612999

Brunoni, A. R., Nitsche, M. A., Bolognini, N., Bikson, M., Wagner, T., Merabet, L., . . . Fregni, F. (2012). Clinical research with transcranial direct current stimulation (tDCS): challenges and future directions. *Brain Stimul, 5*(3), 175-195. doi:10.1016/j.brs.2011.03.002

Brunoni, A. R., Valiengo, L., Baccaro, A., Zanao, T. A., de Oliveira, J. F., Goulart, A., . . . Fregni, F. (2013). The sertraline vs. electrical current therapy for treating depression clinical study: results from a factorial, randomized, controlled trial. *JAMA Psychiatry, 70*(4), 383-391. doi:10.1001/2013.jamapsychiatry.32

Brunoni, A. R., Zanao, T. A., Vanderhasselt, M. A., Valiengo, L., de Oliveira, J. F., Boggio, P. S., . . . Fregni, F. (2014). Enhancement of affective processing induced by bifrontal transcranial direct current stimulation in patients with major depression. *Neuromodulation, 17*(2), 138-142. doi:10.1111/ner.12080

Carleton, R. N., Norton, M. A., & Asmundson, G. J. (2007). Fearing the unknown: a short version of the Intolerance of Uncertainty Scale. *J Anxiety Disord, 21*(1), 105-117. doi:10.1016/j.janxdis.2006.03.014

Cash, R. F. H., Zalesky, A., Thomson, R. H., Tian, Y., Cocchi, L., & Fitzgerald, P. B. (2019). Subgenual Functional Connectivity Predicts Antidepressant Treatment Response to Transcranial Magnetic Stimulation: Independent Validation and Evaluation of Personalization. *Biol Psychiatry, 86*(2), e5-e7. doi:10.1016/j.biopsych.2018.12.002

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. doi:10.1007/s41465-017-0030-7

Cavanagh, J. F., Bismark, A. W., Frank, M. J., & Allen, J. J. B. (2019). Multiple Dissociations Between Comorbid Depression and Anxiety on Reward and Punishment Processing: Evidence From Computationally Informed EEG. *Comput Psychiatr, 3*, 1-17. doi:10.1162/cpsy_a_00024

Chen, N. T. M., Basanovic, J., Notebaert, L., MacLeod, C., & Clarke, P. J. F. (2017). Attentional bias mediates the effect of neurostimulation on emotional vulnerability. *J Psychiatr Res, 93*, 12-19. doi:10.1016/j.jpsychires.2017.05.008

Chen, R., Classen, J., Gerloff, C., Celnik, P., Wassermann, E., Hallett, M., & Cohen, L. G. (1997). Depression of motor cortex excitability by low-frequency transcranial magnetic stimulation. *Neurology, 48*(5), 1398-1403.

Clark, D. A., Steer, R. A., & Beck, A. T. (1994). Common and specific dimensions of self-reported anxiety and depression: implications for the cognitive and tripartite models. *Journal of Abnormal Psychology, 103*(4), 645.

Clarke, P. J., 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. *Biol Psychiatry, 76*(12), 946-952. doi:10.1016/j.biopsych.2014.03.003

Clarke, P. J. F., Van Bockstaele, B., Marinovic, W., Howell, J. A., Boyes, M. E., & Notebaert, L. (2020). The effects of left DLPFC tDCS on emotion regulation, biased attention, and emotional reactivity to negative content. *Cogn Affect Behav Neurosci, 20*(6), 1323-1335. doi:10.3758/s13415-020-00840-2

Cole, E. J., Phillips, A. L., Bentzley, B. S., Stimpson, K. H., Nejad, R., Barmak, F., . . . Williams, N. R. (2022). Stanford Neuromodulation Therapy (SNT): A Double-Blind Randomized Controlled Trial. *Am J Psychiatry, 179*(2), 132-141. doi:10.1176/appi.ajp.2021.20101429

Cooke, S. F., & Bliss, T. V. (2006). Plasticity in the human central nervous system. *Brain, 129*(Pt 7), 1659-1673. doi:10.1093/brain/awl082

Crump, M. J., McDonnell, J. V., & Gureckis, T. M. (2013). Evaluating Amazon's Mechanical Turk as a tool for experimental behavioral research. *PLoS One, 8*(3), e57410. doi:10.1371/journal.pone.0057410

Cyders, M. A., & Coskunpinar, A. (2011). Measurement of constructs using self-report and behavioral lab tasks: is there overlap in nomothetic span and construct representation for impulsivity? *Clin Psychol Rev, 31*(6), 965-982. doi:10.1016/j.cpr.2011.06.001

Dai, Q., Hu, L., & Feng, Z. (2019). Attentional bias modification reduces clinical depression and enhances attention toward happiness. *J Psychiatr Res, 109*, 145-155. doi:10.1016/j.jpsychires.2018.11.024

Daniel-Watanabe, L., McLaughlin, M., Gormley, S., & Robinson, O. J. (2020). Association Between a Directly Translated Cognitive Measure of Negative Bias and Self-reported Psychiatric Symptoms. *Biol Psychiatry Cogn Neurosci Neuroimaging*. doi:10.1016/j.bpsc.2020.02.010

de Aguiar, V., Paolazzi, C. L., & Miceli, G. (2015). tDCS in post-stroke aphasia: the role of stimulation parameters, behavioral treatment and patient characteristics. *Cortex, 63*, 296-316. doi:10.1016/j.cortex.2014.08.015

Dinga, R., Schmaal, L., & Marquand, A. F. (2020). A Closer Look at Depression Biotypes: Correspondence Relating to Grosenick et al. (2019). *Biological Psychiatry: Cognitive Neuroscience and Neuroimaging*. doi:10.1016/j.bpsc.2019.09.011

Dinga, R., Schmaal, L., Penninx, B., van Tol, M. J., Veltman, D. J., van Velzen, L., . . . Marquand, A. F. (2019). Evaluating the evidence for biotypes of depression: Methodological replication and extension of. *Neuroimage Clin, 22*, 101796. doi:10.1016/j.nicl.2019.101796

Disner, S. G., Beevers, C. G., Haigh, E. A., & Beck, A. T. (2011). Neural mechanisms of the cognitive model of depression. *Nat Rev Neurosci, 12*(8), 467-477. doi:10.1038/nrn3027

Dolan, R. J., & Dayan, P. (2013). Goals and habits in the brain. *Neuron, 80*(2), 312-325. doi:10.1016/j.neuron.2013.09.007

Downar, J., Geraci, J., Salomons, T. V., Dunlop, K., Wheeler, S., McAndrews, M. P., . . . Giacobbe, P. (2014). Anhedonia and reward-circuit connectivity distinguish nonresponders from responders to dorsomedial prefrontal repetitive transcranial magnetic stimulation in major depression. *Biol Psychiatry, 76*(3), 176-185. doi:10.1016/j.biopsych.2013.10.026

Drysdale, A. T., Grosenick, L., Downar, J., Dunlop, K., Mansouri, F., Meng, Y., . . . Liston, C. (2017). Resting-state connectivity biomarkers define neurophysiological subtypes of depression. *Nat Med, 23*(1), 28-38. doi:10.1038/nm.4246

Eisenberg, I. W., Bissett, P. G., Zeynep Enkavi, A., Li, J., MacKinnon, D. P., Marsch, L. A., & Poldrack, R. A. (2019). Uncovering the structure of self-regulation through data-driven ontology discovery. *Nat Commun, 10*(1), 2319. doi:10.1038/s41467-019-10301-1

Epp, A. M., Dobson, K. S., Dozois, D. J., & Frewen, P. A. (2012). A systematic meta-analysis of the Stroop task in depression. *Clinical psychology review, 32*(4), 316-328.

Eshel, N., Keller, C. J., Wu, W., Jiang, J., Mills-Finnerty, C., Huemer, J., . . . Etkin, A. (2020). Global connectivity and local excitability changes underlie antidepressant effects of repetitive transcranial magnetic stimulation. *Neuropsychopharmacology*. doi:10.1038/s41386-020-0633-z

Eshel, N., & Roiser, J. P. (2010). Reward and punishment processing in depression. *Biol Psychiatry, 68*(2), 118-124. doi:10.1016/j.biopsych.2010.01.027

Everaert, J., Podina, I. R., & Koster, E. H. W. (2017). A comprehensive meta-analysis of interpretation biases in depression. *Clin Psychol Rev, 58*, 33-48. doi:10.1016/j.cpr.2017.09.005

Farashahi, S., Donahue, C. H., Hayden, B. Y., Lee, D., & Soltani, A. (2019). Flexible combination of reward information across primates. *Nat Hum Behav, 3*(11), 1215-1224. doi:10.1038/s41562-019-0714-3

First, M. B. (2014). Structured clinical interview for the DSM (SCID). *The encyclopedia of clinical psychology*, 1-6.

Foa, E. B., Huppert, J. D., Leiberg, S., Langner, R., Kichic, R., Hajcak, G., & Salkovskis, P. M. (2002). The Obsessive-Compulsive Inventory: Development and validation of a short version. *Psychological Assessment, 14*(4), 485-496. doi:10.1037/1040-3590.14.4.485

Fodor, L. A., Georgescu, R., Cuijpers, P., Szamoskozi, S., David, D., Furukawa, T. A., & Cristea, I. A. (2020). Efficacy of cognitive bias modification interventions in anxiety and depressive disorders: a systematic review and network meta-analysis. *Lancet Psychiatry, 7*(6), 506-514. doi:10.1016/S2215-0366(20)30130-9

Fox, M. D., Buckner, R. L., White, M. P., Greicius, M. D., & Pascual-Leone, A. (2012). Efficacy of transcranial magnetic stimulation targets for depression is related to intrinsic functional connectivity with the subgenual cingulate. *Biol Psychiatry, 72*(7), 595-603. doi:10.1016/j.biopsych.2012.04.028

Fox, M. D., & Greicius, M. (2010). Clinical applications of resting state functional connectivity. *Front Syst Neurosci, 4*, 19. doi:10.3389/fnsys.2010.00019

Fregni, F., Boggio, P. S., Nitsche, M. A., Marcolin, M. A., Rigonatti, S. P., & Pascual-Leone, A. (2006). Treatment of major depression with transcranial direct current stimulation.

Fresco, D. M., Moore, M. T., van Dulmen, M. H., Segal, Z. V., Ma, S. H., Teasdale, J. D., & Williams, J. M. (2007). Initial psychometric properties of the experiences questionnaire: validation of a self-report measure of decentering. *Behav Ther, 38*(3), 234-246. doi:10.1016/j.beth.2006.08.003

Friehs, M. A., & Frings, C. (2019). Offline beats online: transcranial direct current stimulation timing influences on working memory. *NeuroReport, 30*(12), 795-799. doi:10.1097/WNR.0000000000001272

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. doi:10.1016/j.neuron.2010.03.035

Gagne, C., Agai, S., Ramiro, C., Dayan, P., & Bishop, S. (2022). Biased belief priors versus biased belief updating: Differential correlates of depression and anxiety. *PLoS Comput Biol, 18*(8), e1010176. doi:10.1371/journal.pcbi.1010176

Gagne, C., Zika, O., Dayan, P., & Bishop, S. J. (2020). Impaired adaptation of learning to contingency volatility in internalizing psychopathology. *Elife, 9*. doi:10.7554/eLife.61387

Galea, J. M., Vazquez, A., Pasricha, N., de Xivry, J. J., & Celnik, P. (2011). Dissociating the roles of the cerebellum and motor cortex during adaptive learning: the motor cortex retains what the cerebellum learns. *Cereb Cortex, 21*(8), 1761-1770. doi:10.1093/cercor/bhq246

Gard, D. E., Gard, M. G., Kring, A. M., & John, O. P. (2006). Anticipatory and consummatory components of the experience of pleasure: A scale development study. *Journal of Research in Personality, 40*(6), 1086-1102. doi:10.1016/j.jrp.2005.11.001

George, M. S., Ketter, T. A., & Post, R. M. (1994). Prefrontal cortex dysfunction in clinical depression. *Depression, 2*(2), 59-72. doi:10.1002/depr.3050020202

George, M. S., Lisanby, S. H., Avery, D., McDonald, W. M., Durkalski, V., Pavlicova, M., . . . Zarkowski, P. (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.
Gillan, C. M., Kalanthroff, E., Evans, M., Weingarden, H. M., Jacoby, R. J., Gershkovich, M., . . . Simpson, H. B. (2020). Comparison of the Association Between Goal-Directed Planning and Self-reported Compulsivity vs Obsessive-Compulsive Disorder Diagnosis. *JAMA Psychiatry, 77*(1), 77-85. doi:10.1001/jamapsychiatry.2019.2998
Gillan, C. M., Kosinski, M., Whelan, R., Phelps, E. A., & Daw, N. D. (2016). Characterizing a psychiatric symptom dimension related to deficits in goal-directed control. *Elife, 5*. doi:10.7554/eLife.11305
Godlewska, B. R., Browning, M., Norbury, R., Cowen, P. J., & Harmer, C. J. (2016). Early changes in emotional processing as a marker of clinical response to SSRI treatment in depression. *Transl Psychiatry, 6*(11), e957. doi:10.1038/tp.2016.130
Godlewska, B. R., Norbury, R., Selvaraj, S., Cowen, P. J., & Harmer, C. J. (2012). Short-term SSRI treatment normalises amygdala hyperactivity in depressed patients. *Psychol Med, 42*(12), 2609-2617. doi:10.1017/S0033291712000591
Gotlib, I. H. (1983). Perception and recall of interpersonal feedback: Negative bias in depression. *Cognitive Therapy and Research, 7*(5), 399-412.
Grahek, I., Shenhav, A., Musslick, S., Krebs, R. M., & Koster, E. H. W. (2019). Motivation and cognitive control in depression. *Neurosci Biobehav Rev, 102*, 371-381. doi:10.1016/j.neubiorev.2019.04.011
Grasso, P. A., Tonolli, E., & Miniussi, C. (2020). Effects of different transcranial direct current stimulation protocols on visuo-spatial contextual learning formation: evidence of homeostatic regulatory mechanisms. *Sci Rep, 10*(1), 4622. doi:10.1038/s41598-020-61626-7
Greicius, M. D., Flores, B. H., Menon, V., Glover, G. H., Solvason, H. B., Kenna, H., . . . Schatzberg, A. F. (2007). Resting-state functional connectivity in major depression: abnormally increased contributions from subgenual cingulate cortex and thalamus. *Biol Psychiatry, 62*(5), 429-437. doi:10.1016/j.biopsych.2006.09.020
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. *Biol Psychiatry, 63*(4), 369-376. doi:10.1016/j.biopsych.2007.05.033
Gros, D. F., Antony, M. M., Simms, L. J., & McCabe, R. E. (2007). Psychometric properties of the State-Trait Inventory for Cognitive and Somatic Anxiety (STICSA): comparison to the State-Trait Anxiety Inventory (STAI). *Psychol Assess, 19*(4), 369-381. doi:10.1037/1040-3590.19.4.369
Gur, R. C., Erwin, R. J., Gur, R. E., Zwil, A. S., Heimberg, C., & Kraemer, H. C. (1992). Facial emotion discrimination: II. Behavioral findings in depression. *Psychiatry research, 42*(3), 241-251.
Haber, S. N., & Knutson, B. (2010). The reward circuit: linking primate anatomy and human imaging. *Neuropsychopharmacology, 35*(1), 4-26. doi:10.1038/npp.2009.129
Hägele, C., Schlagenhauf, F., Rapp, M., Sterzer, P., Beck, A., Bermpohl, F., . . . Heinz, A. (2015). Dimensional psychiatry: reward dysfunction and depressive mood across psychiatric disorders. *Psychopharmacology (Berl), 232*(2), 331-341. doi:10.1007/s00213-014-3662-7

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. *Sci Rep, 6*, 25160. doi:10.1038/srep25160

Harle, K. M., Guo, D., Zhang, S., Paulus, M. P., & Yu, A. J. (2017). Anhedonia and anxiety underlying depressive symptomatology have distinct effects on reward-based decision-making. *PLoS One, 12*(10), e0186473. doi:10.1371/journal.pone.0186473

Harmer, C. J., Bhagwagar, Z., Perrett, D. I., Vollm, 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. doi:10.1038/sj.npp.1300004

Harmer, C. J., Goodwin, G. M., & Cowen, P. J. (2009). Why do antidepressants take so long to work? A cognitive neuropsychological model of antidepressant drug action. *Br J Psychiatry, 195*(2), 102-108. doi:10.1192/bjp.bp.108.051193

Hecht, D. (2010). Depression and the hyperactive right-hemisphere. *Neurosci Res, 68*(2), 77-87. doi:10.1016/j.neures.2010.06.013

Heeger, D., & Landy, M. (1997). Signal detection theory. *Dept. Psych., Stanford Univ., Stanford, CA, Teaching Handout*.

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), e0124182. doi:10.1371/journal.pone.0124182

Heinz, A., Schlagenhauf, F., Beck, A., & Wackerhagen, C. (2016). Dimensional psychiatry: mental disorders as dysfunctions of basic learning mechanisms. *J Neural Transm (Vienna), 123*(8), 809-821. doi:10.1007/s00702-016-1561-2

Herrington, J. D., Mohanty, A., Koven, N. S., Fisher, J. E., Stewart, J. L., Banich, M. T., . . . Heller, W. (2005). Emotion-modulated performance and activity in left dorsolateral prefrontal cortex. *Emotion, 5*(2), 200-207. doi:10.1037/1528-3542.5.2.200

Herzfeld, D. J., Pastor, D., Haith, A. M., Rossetti, Y., Shadmehr, R., & O'Shea, J. (2014). Contributions of the cerebellum and the motor cortex to acquisition and retention of motor memories. *Neuroimage, 98*, 147-158. doi:10.1016/j.neuroimage.2014.04.076

Horne, S. J., Topp, T. E., & Quigley, L. (2021). Depression and the willingness to expend cognitive and physical effort for rewards: A systematic review. *Clin Psychol Rev, 88*, 102065. doi:10.1016/j.cpr.2021.102065

Hotelling, H. (1936). Relations Between Two Sets of Variates. *Biometrika, 28*(3/4), 321-377. doi:10.2307/2333955

Hoven, M., Denys, D., Rouault, M., Luigjes, J., & van Holst, R. (2022). How do confidence and self-beliefs relate in psychopathology: a transdiagnostic approach. *PsyArXiv. March, 11*.

Huang, C., & Chen, J. H. (2015). Meta-Analysis of the Factor Structures of the Beck Depression Inventory-II. *Assessment, 22*(4), 459-472. doi:10.1177/1073191114548873

Huang, H., Thompson, W., & Paulus, M. P. (2017). Computational Dysfunctions in Anxiety: Failure to Differentiate Signal From Noise. *Biol Psychiatry, 82*(6), 440-446. doi:10.1016/j.biopsych.2017.07.007

Huang, Y. Z., Lu, M. K., Antal, A., Classen, J., Nitsche, M., Ziemann, U., . . . Rothwell, J. (2017). Plasticity induced by non-invasive transcranial brain stimulation: A position paper. *Clin Neurophysiol, 128*(11), 2318-2329. doi:10.1016/j.clinph.2017.09.007

Huys, Q. J. M., 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. doi:10.1186/2045-5380-3-12

Insel, T., Cuthbert, B., Garvey, M., Heinssen, R., Pine, D. S., Quinn, K., . . . Wang, P. (2010). Research domain criteria (RDoC): toward a new classification framework for research on mental disorders. In (Vol. 167, pp. 748-751): Am Psychiatric Assoc.

Ironside, M., Browning, M., Ansari, T. L., Harvey, C. J., Sekyi-Djan, M. N., Bishop, S. J., . . . O'Shea, J. (2019). Effect of Prefrontal Cortex Stimulation on Regulation of Amygdala Response to Threat in Individuals With Trait Anxiety: A Randomized Clinical Trial. *JAMA Psychiatry, 76*(1), 71-78. doi:10.1001/jamapsychiatry.2018.2172

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. *Biol Psychiatry, 79*(10), 823-830. doi:10.1016/j.biopsych.2015.06.012

Kammer, T., Beck, S., Thielscher, A., Laubis-Herrmann, U., & Topka, H. (2001). Motor thresholds in humans: a transcranial magnetic stimulation study comparing different pulse waveforms, current directions and stimulator types. *Clinical Neurophysiology, 112*(2), 250-258.

Karabanov, A. N., Saturnino, G. B., Thielscher, A., & Siebner, H. R. (2019). Can Transcranial Electrical Stimulation Localize Brain Function? *Front Psychol, 10*, 213. doi:10.3389/fpsyg.2019.00213

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. *J Neurosci, 31*(43), 15284-15293. doi:10.1523/JNEUROSCI.0542-11.2011

Kolling, N., Scholl, J., Chekroud, A., Trier, H. A., & Rushworth, M. F. S. (2018). Prospection, Perseverance, and Insight in Sequential Behavior. *Neuron, 99*(5), 1069-1082 e1067. doi:10.1016/j.neuron.2018.08.018

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. *Psychol Med, 44*(3), 579-592. doi:10.1017/S0033291713001074

Kube, T., Schwarting, R., Rozenkrantz, L., Glombiewski, J. A., & Rief, W. (2020). Distorted Cognitive Processes in Major Depression: A Predictive Processing Perspective. *Biol Psychiatry, 87*(5), 388-398. doi:10.1016/j.biopsych.2019.07.017

Kumar, P., Goer, F., Murray, L., Dillon, D. G., Beltzer, M. L., Cohen, A. L., . . . Pizzagalli, D. A. (2018). Impaired reward prediction error encoding and striatal-midbrain connectivity in depression. *Neuropsychopharmacology, 43*(7), 1581-1588. doi:10.1038/s41386-018-0032-x

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. *J Behav Ther Exp Psychiatry, 43*(4), 1088-1094. doi:10.1016/j.jbtep.2012.05.007

Lee, J. H., Lee, T. L., & Kang, N. (2021). Transcranial Direct Current Stimulation Decreased Cognition-Related Reaction Time in Older Adults: A Systematic Review and Meta-Analysis. *Ageing Res Rev*, 101377. doi:10.1016/j.arr.2021.101377

Lefaucheur, J. P., Andre-Obadia, N., Antal, A., Ayache, S. S., Baeken, C., Benninger, D. H., . . . Garcia-Larrea, L. (2014). Evidence-based guidelines on the therapeutic use of repetitive transcranial magnetic stimulation (rTMS). *Clin Neurophysiol, 125*(11), 2150-2206. doi:10.1016/j.clinph.2014.05.021

Lepping, P., Schonfeldt-Lecuona, C., Sambhi, R. S., Lanka, S. V., Lane, S., Whittington, R., . . . Poole, R. (2014). A systematic review of the clinical relevance of repetitive transcranial magnetic stimulation. *Acta Psychiatr Scand, 130*(5), 326-341. doi:10.1111/acps.12276

Levkovitz, Y., Isserles, M., Padberg, F., Lisanby, S. H., Bystritsky, A., Xia, G., . . . Dannon, P. (2015). Efficacy and safety of deep transcranial magnetic stimulation for major depression: a prospective multicenter randomized controlled trial. *World Psychiatry, 14*(1), 64-73.

Li, L. M., Violante, I. R., Leech, R., Ross, E., Hampshire, A., Opitz, A., . . . Sharp, D. J. (2019). Brain state and polarity dependent modulation of brain networks by transcranial direct current stimulation. *Hum Brain Mapp, 40*(3), 904-915. doi:10.1002/hbm.24420

Lim, G. Y., Tam, W. W., Lu, Y., Ho, C. S., Zhang, M. W., & Ho, R. C. (2018). Prevalence of Depression in the Community from 30 Countries between 1994 and 2014. *Sci Rep, 8*(1), 2861. doi:10.1038/s41598-018-21243-x

Liston, C., Chen, A. C., Zebley, B. D., Drysdale, A. T., Gordon, R., Leuchter, B., . . . Dubin, M. J. (2014). Default mode network mechanisms of transcranial magnetic stimulation in depression. *Biol Psychiatry, 76*(7), 517-526. doi:10.1016/j.biopsych.2014.01.023

Liu, A., Voroslakos, M., Kronberg, G., Henin, S., Krause, M. R., Huang, Y., . . . Buzsaki, G. (2018). Immediate neurophysiological effects of transcranial electrical stimulation. *Nat Commun, 9*(1), 5092. doi:10.1038/s41467-018-07233-7

Loo, C. K., Sachdev, P. S., Haindl, W., Wen, W., Mitchell, P. B., Croker, V. M., & Malhi, G. S. (2003). High (15 Hz) and low (1 Hz) frequency transcranial magnetic stimulation have different acute effects on regional cerebral blood flow in depressed patients. *Psychol Med, 33*(6), 997-1006. doi:10.1017/s0033291703007955

MacLeod, C., Rutherford, E., Campbell, L., Ebsworthy, G., & Holker, L. (2002). Selective attention and emotional vulnerability: Assessing the causal basis of their association through the experimental manipulation of attentional bias. *Journal of Abnormal Psychology, 111*(1), 107-123. doi:10.1037/0021-843x.111.1.107

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. *Exp Brain Res, 232*(10), 3345-3351. doi:10.1007/s00221-014-4022-x

Mason, O., Linney, Y., & Claridge, G. (2005). Short scales for measuring schizotypy. *Schizophr Res, 78*(2-3), 293-296. doi:10.1016/j.schres.2005.06.020

Massi, B., Donahue, C. H., & Lee, D. (2018). Volatility Facilitates Value Updating in the Prefrontal Cortex. *Neuron, 99*(3), 598-608 e594. doi:10.1016/j.neuron.2018.06.033

Mayberg, H. S., Lozano, A. M., Voon, V., McNeely, H. E., Seminowicz, D., Hamani, C., . . . Kennedy, S. H. (2005). Deep brain stimulation for treatment-resistant depression. *Neuron, 45*(5), 651-660. doi:10.1016/j.neuron.2005.02.014

McHugh, R. K., Daughters, S. B., Lejuez, C. W., Murray, H. W., Hearon, B. A., Gorka, S. M., & Otto, M. W. (2011). Shared Variance among Self-Report and Behavioral Measures of Distress Intolerance. *Cognit Ther Res, 35*(3), 266-275. doi:10.1007/s10608-010-9295-1

Menon, V. (2011). Large-scale brain networks and psychopathology: a unifying triple network model. *Trends Cogn Sci, 15*(10), 483-506. doi:10.1016/j.tics.2011.08.003

Mihalik, A., Adams, R. A., & Huys, Q. (2020). Canonical Correlation Analysis for Identifying Biotypes of Depression. *Biol Psychiatry Cogn Neurosci Neuroimaging, 5*(5), 478-480. doi:10.1016/j.bpsc.2020.02.002

Mukherjee, D., Filipowicz, A. L. S., Vo, K., Satterthwaite, T. D., & Kable, J. W. (2020). Reward and punishment reversal-learning in major depressive disorder. *J Abnorm Psychol, 129*(8), 810-823. doi:10.1037/abn0000641

Müller-Dahlhaus, F., & Ziemann, U. (2015). Metaplasticity in human cortex. *Neuroscientist, 21*(2), 185-202. doi:10.1177/1073858414526645

Mundt, J. C., Marks, I. M., Shear, M. K., & Greist, J. H. (2002). The Work and Social Adjustment Scale: a simple measure of impairment in functioning. *Br J Psychiatry, 180*, 461-464. doi:10.1192/bjp.180.5.461

National Institute for Clinical Excellence. (2015a, August 2020). Repetitive transcranial magnetic stimulation for depression. Retrieved from https://www.nice.org.uk/guidance/ipg542

National Institute for Clinical Excellence. (2015b). Transcranial direct current stimulation (tDCS) for depression. Retrieved from http://nice.org.uk/guidance/ipg530

Nejati, V., Khalaji, S., Goodarzi, H., & Nitsche, M. (2021). The role of ventromedial and dorsolateral prefrontal cortex in attention and interpretation biases in individuals with general anxiety disorder (GAD): A tDCS study. *Journal of Psychiatric Research, 144*, 269-277. doi:10.1016/j.jpsychires.2021.10.034

Nitsche, M. A., Koschack, J., Pohlers, H., Hullemann, S., Paulus, W., & Happe, S. (2012). Effects of frontal transcranial direct current stimulation on emotional state and processing in healthy humans. *Front Psychiatry, 3*, 58. doi:10.3389/fpsyt.2012.00058

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*(Pt 3), 633.

O'Reardon, J. P., Solvason, H. B., Janicak, P. G., Sampson, S., Isenberg, K. E., Nahas, Z., . . . Sackeim, H. A. (2007). Efficacy and safety of transcranial magnetic stimulation in the acute treatment of major depression: a multisite randomized controlled trial. *Biol Psychiatry, 62*(11), 1208-1216. doi:10.1016/j.biopsych.2007.01.018

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*. doi:10.7554/eLife.26602

Oldrati, V., Colombo, B., & Antonietti, A. (2018). Combination of a short cognitive training and tDCS to enhance visuospatial skills: A comparison between online and offline neuromodulation. *Brain Res, 1678*, 32-39. doi:10.1016/j.brainres.2017.10.002

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. doi:10.1016/j.neuroimage.2015.01.033

Overman, M. J., Browning, M., & O'Shea, J. (2020). Inducing Affective Learning Biases with Cognitive Training and Prefrontal tDCS: A Proof-of-Concept Study. *Cognitive Therapy and Research*. doi:10.1007/s10608-020-10146-9

Overman, M. J., Sarrazin, V., Browning, M., & O'Shea, J. (2021). Stimulating human prefrontal cortex increases reward learning. *bioRxiv*.

Panitz, M., Deserno, L., Kaminski, E., Villringer, A., Sehm, B., & Schlagenhauf, F. (2022). Anodal tDCS over the medial prefrontal cortex enhances behavioral adaptation after punishments during reversal learning through increased updating of unchosen choice options. *Cereb Cortex Commun, 3*(1), tgac006. doi:10.1093/texcom/tgac006

Pantazatos, S. P., McLntosh, J. R., Saber, G. T., Sun, X., Doose, J., Faller, J., . . . Brown, T. R. (2022). Functional and effective connectivity between dorsolateral prefrontal and subgenual anterior cingulate cortex depends on the timing of transcranial magnetic stimulation relative to the phase of prefrontal alpha EEG. *bioRxiv*. doi:10.1101/2022.02.14.480466

Pascual-Leone, A., Rubio, B., Pallardó, F., & Catalá, M. D. (1996). Rapid-rate transcranial magnetic stimulation of left dorsolateral prefrontal cortex in drug-resistant depression. *The Lancet, 348*(9022), 233-237. doi:10.1016/s0140-6736(96)01219-6

Pascual-Leone, A., Valls-Solé, J., Wassermann, E. M., & Hallett, M. (1994). Responses to rapid-rate transcranial magnetic stimulation of the human motor cortex. *Brain, 117*(4), 847-858.

Pena-Gomez, C., Sala-Lonch, R., Junque, C., Clemente, I. C., Vidal, D., Bargallo, N., . . . Bartres-Faz, D. (2012). Modulation of large-scale brain networks by transcranial direct current stimulation evidenced by resting-state functional MRI. *Brain Stimul, 5*(3), 252-263. doi:10.1016/j.brs.2011.08.006

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 Stimul, 9*(3), 336-346. doi:10.1016/j.brs.2016.03.010

Pessiglione, M., Seymour, B., Flandin, G., Dolan, R. J., & Frith, C. D. (2006). Dopamine-dependent prediction errors underpin reward-seeking behaviour in humans. *Nature, 442*(7106), 1042-1045. doi:10.1038/nature05051

Philip, N. S., Barredo, J., van 't Wout-Frank, M., Tyrka, A. R., Price, L. H., & Carpenter, L. L. (2018). Network Mechanisms of Clinical Response to Transcranial Magnetic Stimulation in Posttraumatic Stress Disorder and Major Depressive Disorder. *Biol Psychiatry, 83*(3), 263-272. doi:10.1016/j.biopsych.2017.07.021

Pike, A. C., & Robinson, O. J. (2022). Reinforcement Learning in Patients With Mood and Anxiety Disorders vs Control Individuals: A Systematic Review and Meta-analysis. *JAMA Psychiatry*. doi:10.1001/jamapsychiatry.2022.0051

Pirulli, C., Fertonani, A., & Miniussi, C. (2013). The role of timing in the induction of neuromodulation in perceptual learning by transcranial electric stimulation. *Brain Stimul, 6*(4), 683-689. doi:10.1016/j.brs.2012.12.005

Pizzagalli, D. A., Holmes, A. J., Dillon, D. G., Goetz, E. L., Birk, J. L., Bogdan, R., . . . Fava, M. (2009). Reduced caudate and nucleus accumbens response to rewards in unmedicated individuals with major depressive disorder. *American Journal of Psychiatry, 166*(6), 702-710.

Pizzagalli, D. A., Jahn, A. L., & O'Shea, J. P. (2005). Toward an objective characterization of an anhedonic phenotype: a signal-detection approach. *Biol Psychiatry, 57*(4), 319-327. doi:10.1016/j.biopsych.2004.11.026

Poydasheva, A. G., Bakulin, I. S., Lagoda, D. Y., Pavlova, E. L., Suponeva, N. A., & Piradov, M. A. (2021). High-Definition Transcranial Direct Current Electrical Stimulation. *Neuroscience and Behavioral Physiology, 51*(8), 1190-1198. doi:10.1007/s11055-021-01178-z

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.

Pulcu, E., & Browning, M. (2017). Affective bias as a rational response to the statistics of rewards and punishments. *Elife, 6*. doi:10.7554/eLife.27879

Pulcu, E., & Browning, M. (2019). The Misestimation of Uncertainty in Affective Disorders. *Trends Cogn Sci, 23*(10), 865-875. doi:10.1016/j.tics.2019.07.007

Pulcu, E., Shkreli, L., Holst, C. G., Woud, M. L., Craske, M. G., Browning, M., & Reinecke, A. (2019). The Effects of the Angiotensin II Receptor Antagonist Losartan on Appetitive Versus Aversive Learning: A Randomized Controlled Trial. *Biol Psychiatry, 86*(5), 397-404. doi:10.1016/j.biopsych.2019.04.010

Pulcu, E., Trotter, P. D., Thomas, E. J., McFarquhar, M., Juhasz, G., Sahakian, B. J., . . . Elliott, R. (2014). Temporal discounting in major depressive disorder. *Psychol Med, 44*(9), 1825-1834. doi:10.1017/S0033291713002584

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. *J Physiol, 591*(10), 2563-2578. doi:10.1113/jphysiol.2012.247171

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. *J Neurosci, 35*(43), 14544-14556. doi:10.1523/JNEUROSCI.2322-15.2015

Razza, L. B., Palumbo, P., Moffa, A. H., Carvalho, A. F., Solmi, M., Loo, C. K., & Brunoni, A. R. (2020). A systematic review and meta-analysis on the effects of transcranial direct current stimulation in depressive episodes. *Depress Anxiety, 37*(7), 594-608. doi:10.1002/da.23004

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. *Proc Natl Acad Sci U S A, 106*(5), 1590-1595. doi:10.1073/pnas.0805413106

Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. *Current research and theory*, 64-99.

Richieri, R., Jouvenoz, D., Verger, A., Fiat, P., Boyer, L., Lançon, C., & Guedj, E. (2017). Changes in dorsolateral prefrontal connectivity after rTMS in treatment-resistant depression: a brain perfusion SPECT study. *Eur J Nucl Med Mol Imaging, 44*(6), 1051-1055. doi:10.1007/s00259-017-3640-5

Rock, P. L., Roiser, J. P., Riedel, W. J., & Blackwell, A. D. (2014). Cognitive impairment in depression: a systematic review and meta-analysis. *Psychol Med, 44*(10), 2029-2040. doi:10.1017/S0033291713002535

Roiser, J. P., Elliott, R., & Sahakian, B. J. (2012). Cognitive mechanisms of treatment in depression. *Neuropsychopharmacology, 37*(1), 117-136. doi:10.1038/npp.2011.183

Rouault, M., Seow, T., Gillan, C. M., & Fleming, S. M. (2018). Psychiatric Symptom Dimensions Are Associated With Dissociable Shifts in Metacognition but Not Task Performance. *Biol Psychiatry, 84*(6), 443-451. doi:10.1016/j.biopsych.2017.12.017

Rush, A. J., Trivedi, M. H., Ibrahim, H. M., Carmody, T. J., Arnow, B., Klein, D. N., . . . Keller, M. B. (2003). The 16-Item quick inventory of depressive symptomatology (QIDS), clinician rating (QIDS-C), and self-report (QIDS-SR): a psychometric evaluation in patients with chronic major depression. *Biological Psychiatry, 54*(5), 573-583. doi:10.1016/s0006-3223(02)01866-8

Rush, A. J., Trivedi, M. H., Wisniewski, S. R., Nierenberg, A. A., Stewart, J. W., Warden, D., . . . Lebowitz, B. D. (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. (2008). Intention, choice, and the medial frontal cortex. *Ann N Y Acad Sci, 1124*, 181-207. doi:10.1196/annals.1440.014

Salomons, T. V., Dunlop, K., Kennedy, S. H., Flint, A., Geraci, J., Giacobbe, P., & Downar, J. (2014). Resting-state cortico-thalamic-striatal connectivity predicts response to dorsomedial prefrontal rTMS in major depressive disorder. *Neuropsychopharmacology, 39*(2), 488-498. doi:10.1038/npp.2013.222

Sanchez-Lopez, A., Vanderhasselt, M. A., Allaert, J., Baeken, C., & De Raedt, R. (2018). Neurocognitive mechanisms behind emotional attention: Inverse effects of anodal tDCS over the left and right DLPFC on gaze disengagement from emotional faces. *Cogn Affect Behav Neurosci, 18*(3), 485-494. doi:10.3758/s13415-018-0582-8

Santomauro, D. F., Mantilla Herrera, A. M., Shadid, J., Zheng, P., Ashbaugh, C., Pigott, D. M., . . . Ferrari, A. J. (2021). Global prevalence and burden of depressive and anxiety disorders in 204 countries and territories in 2020 due to the COVID-19 pandemic. *The Lancet, 398*(10312), 1700-1712. doi:10.1016/s0140-6736(21)02143-7

Sarrazin, V., & O'Shea, J. (2021). The Value of Neuroimaging for Treating Depression with Brain Stimulation. *Transcranial Direct Current Stimulation in Neuropsychiatric Disorders*, 173-210.

Scholl, J., & Klein-Flugge, M. (2018). Understanding psychiatric disorder by capturing ecologically relevant features of learning and decision-making. *Behav Brain Res, 355*, 56-75. doi:10.1016/j.bbr.2017.09.050

Scholl, J., Trier, H. A., Rushworth, M. F. S., & Kolling, N. (2022). The effect of apathy and compulsivity on planning and stopping in sequential decision-making. *PLoS Biol, 20*(3), e3001566. doi:10.1371/journal.pbio.3001566

Seow, T. X. F., & Gillan, C. M. (2020). Transdiagnostic Phenotyping Reveals a Host of Metacognitive Deficits Implicated in Compulsivity. *Sci Rep, 10*(1), 2883. doi:10.1038/s41598-020-59646-4

Sharma, M., Farahani, F., Bikson, M., & Parra, L. C. (2021). Animal Studies on the Mechanisms of Low-Intensity Transcranial Electric Stimulation. In *Transcranial Direct Current Stimulation in Neuropsychiatric Disorders* (pp. 67-92): Springer.

Sharot, T., Korn, C. W., & Dolan, R. J. (2011). How unrealistic optimism is maintained in the face of reality. *Nat Neurosci, 14*(11), 1475-1479. doi:10.1038/nn.2949

Shaw, M. T., Kasschau, M., Dobbs, B., Pawlak, N., Pau, W., Sherman, K., . . . Charvet, L. E. (2017). Remotely Supervised Transcranial Direct Current Stimulation: An Update on Safety and Tolerability. *J Vis Exp*(128). doi:10.3791/56211

Siddiqi, S. H., Schaper, F., Horn, A., Hsu, J., Padmanabhan, J. L., Brodtmann, A., . . . Fox, M. D. (2021). Brain stimulation and brain lesions converge on common causal circuits in neuropsychiatric disease. *Nat Hum Behav*. doi:10.1038/s41562-021-01161-1

Siddiqi, S. H., Taylor, S. F., Cooke, D., Pascual-Leone, A., George, M. S., & Fox, M. D. (2020a). Distinct Symptom-Specific Treatment Targets for Circuit-Based Neuromodulation. *Am J Psychiatry*, appiajp201919090915. doi:10.1176/appi.ajp.2019.19090915

Siddiqi, S. H., Taylor, S. F., Cooke, D., Pascual-Leone, A., George, M. S., & Fox, M. D. (2020b). Distinct Symptom-Specific Treatment Targets for Circuit-Based Neuromodulation. *Am J Psychiatry, 177*(5), 435-446. doi:10.1176/appi.ajp.2019.19090915

Siddiqi, S. H., Weigand, A., Cooke, D., Pascual-Leone, A., & Fox, M. D. (2019). Abstract #155: Individualized connectivity between rTMS targets and the subgenual cingulate is unrelated to antidepressant response. *Brain Stimulation, 12*(2). doi:10.1016/j.brs.2018.12.162

Silvanto, J., Muggleton, N., & Walsh, V. (2008). State-dependency in brain stimulation studies of perception and cognition. *Trends Cogn Sci, 12*(12), 447-454. doi:10.1016/j.tics.2008.09.004

Smith, S. M., Nichols, T. E., Vidaurre, D., Winkler, A. M., Behrens, T. E., Glasser, M. F., . . . Miller, K. L. (2015). A positive-negative mode of population covariation links brain connectivity, demographics and behavior. *Nat Neurosci, 18*(11), 1565-1567. doi:10.1038/nn.4125

Snaith, R., Hamilton, M., Morley, S., Humayan, A., Hargreaves, D., & Trigwell, P. (1995). A scale for the assessment of hedonic tone the Snaith–Hamilton Pleasure Scale. *The British Journal of Psychiatry, 167*(1), 99-103.

Sommer, M., Alfaro, A., Rummel, M., Speck, S., Lang, N., Tings, T., & Paulus, W. (2006). Half sine, monophasic and biphasic transcranial magnetic stimulation of the human motor cortex. *Clin Neurophysiol, 117*(4), 838-844. doi:10.1016/j.clinph.2005.10.029

Speer, A. M., Kimbrell, T. A., Wassermann, E. M., J, D. R., Willis, M. W., Herscovitch, P., & Post, R. M. (2000). Opposite effects of high and low frequency rTMS on regional brain activity in depressed patients. *Biol Psychiatry, 48*(12), 1133-1141. doi:10.1016/s0006-3223(00)01065-9

Spielberger, C. (1983). Manual for the State-Trait Anxiety Inventory; Palo Alto, CA, Ed. In: Consulting Psychologists Press, Inc.: Columbia, MO, USA.

Spronk, D., Arns, M., & Fitzgerald, P. B. (2011). Repetitive Transcranial Magnetic Stimulation in Depression. In *Neurofeedback and Neuromodulation Techniques and Applications* (pp. 257-291).

Stagg, C. J., Jayaram, G., Pastor, D., Kincses, Z. T., Matthews, P. M., & Johansen-Berg, H. (2011). Polarity and timing-dependent effects of transcranial direct current stimulation in explicit motor learning. *Neuropsychologia, 49*(5), 800-804. doi:10.1016/j.neuropsychologia.2011.02.009

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. *J Neurosci, 33*(28), 11425-11431. doi:10.1523/JNEUROSCI.3887-12.2013

Stan Developement Team. (2022). Stan Modeling Language Users Guide and Reference Manual, version 2.30. doi:https://mc-stan.org

Steer, R. A., Clark, D. A., Beck, A. T., & Ranieri, W. F. (1995). Common and specific dimensions of self-reported anxiety and depression: a replication. *Journal of Abnormal Psychology, 104*(3), 542.

Summers, J. J., Kang, N., & Cauraugh, J. H. (2016). Does transcranial direct current stimulation enhance cognitive and motor functions in the ageing brain? A systematic review and meta- analysis. *Ageing Res Rev, 25*, 42-54. doi:10.1016/j.arr.2015.11.004

Taylor, S. F., Ho, S. S., Abagis, T., Angstadt, M., Maixner, D. F., Welsh, R. C., & Hernandez-Garcia, L. (2018). Changes in brain connectivity during a sham-controlled, transcranial magnetic stimulation trial for depression. *J Affect Disord, 232*, 143-151. doi:10.1016/j.jad.2018.02.019

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. doi:10.1016/j.neuroimage.2008.05.049

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. doi:10.1016/j.cortex.2014.08.026

Uher, R., Dernovsek, M. Z., Mors, O., Hauser, J., Souery, D., Zobel, A., . . . Farmer, A. (2011). Melancholic, atypical and anxious depression subtypes and outcome of treatment with escitalopram and nortriptyline. *J Affect Disord, 132*(1-2), 112-120. doi:10.1016/j.jad.2011.02.014

Vanderhasselt, M. A., De Raedt, R., Brunoni, A. R., Campanha, 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), e62219. doi:10.1371/journal.pone.0062219

Vehtari, A., Gelman, A., Simpson, D., Carpenter, B., & Bürkner, P.-C. (2021). Rank-Normalization, Folding, and Localization: An Improved Rˆ for Assessing Convergence of MCMC (with Discussion). *Bayesian Analysis, 16*(2). doi:10.1214/20-ba1221

Voineskos, D., Daskalakis, Z. J., & Blumberger, D. M. (2020). Management of Treatment-Resistant Depression: Challenges and Strategies. *Neuropsychiatr Dis Treat, 16*, 221-234. doi:10.2147/NDT.S198774

Voon, V., Reiter, A., Sebold, M., & Groman, S. (2017). Model-Based Control in Dimensional Psychiatry. *Biol Psychiatry, 82*(6), 391-400. doi:10.1016/j.biopsych.2017.04.006

Vos, T., Lim, S. S., Abbafati, C., Abbas, K. M., Abbasi, M., Abbasifard, M., . . . Murray, C. J. L. (2020). Global burden of 369 diseases and injuries in 204 countries and territories, 1990–2019: a systematic analysis for the Global Burden of Disease Study 2019. *The Lancet, 396*(10258), 1204-1222. doi:10.1016/s0140-6736(20)30925-9

Walsh, A. E. L., Browning, M., Drevets, W. C., Furey, M., & Harmer, C. J. (2018). Dissociable temporal effects of bupropion on behavioural measures of emotional and reward processing in depression. *Philos Trans R Soc Lond B Biol Sci, 373*(1742). doi:10.1098/rstb.2017.0030

Walsh, A. E. L., Huneke, N. T. M., Brown, R., Browning, M., Cowen, P., & Harmer, C. J. (2018). A Dissociation of the Acute Effects of Bupropion on Positive Emotional Processing and Reward Processing in Healthy Volunteers. *Front Psychiatry, 9*, 482. doi:10.3389/fpsyt.2018.00482

Wang, H. T., Smallwood, J., Mourao-Miranda, J., Xia, C. H., Satterthwaite, T. D., Bassett, D. S., & Bzdok, D. (2020). Finding the needle in a high-dimensional haystack: Canonical correlation analysis for neuroscientists. *Neuroimage, 216*, 116745. doi:10.1016/j.neuroimage.2020.116745

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.

Westwood, S. J., & Romani, C. (2017). Transcranial direct current stimulation (tDCS) modulation of picture naming and word reading: A meta-analysis of single session tDCS applied to healthy participants. *Neuropsychologia, 104*, 234-249. doi:10.1016/j.neuropsychologia.2017.07.031

Wiethoff, S., Hamada, M., & Rothwell, J. C. (2014). Variability in response to transcranial direct current stimulation of the motor cortex. *Brain Stimul, 7*(3), 468-475. doi:10.1016/j.brs.2014.02.003

Winkler, A. M., Renaud, O., Smith, S. M., & Nichols, T. E. (2020). Permutation inference for canonical correlation analysis. *Neuroimage, 220*, 117065. doi:10.1016/j.neuroimage.2020.117065

Wyczesany, M., Capotosto, P., Zappasodi, F., & Prete, G. (2018). Hemispheric asymmetries and emotions: Evidence from effective connectivity. *Neuropsychologia, 121*, 98-105. doi:10.1016/j.neuropsychologia.2018.10.007

Young, A. W., Rowland, D., Calder, A. J., Etcoff, N. L., Seth, A., & Perrett, D. I. (1997). Facial expression megamix: Tests of dimensional and category accounts of emotion recognition. *Cognition, 63*(3), 271-313.

Zhang, B., Liu, J., Bao, T., Wilson, G., Park, J., Zhao, B., & Kong, J. (2020). Locations for noninvasive brain stimulation in treating depressive disorders: A combination of meta-analysis and resting-state functional connectivity analysis. *Aust N Z J Psychiatry, 54*(6), 582-590. doi:10.1177/0004867420920372

Ziemann, U., Paulus, W., Nitsche, M. A., Pascual-Leone, A., Byblow, W. D., Berardelli, A., . . . Rothwell, J. C. (2008). Consensus: Motor cortex plasticity protocols. *Brain Stimul, 1*(3), 164-182. doi:10.1016/j.brs.2008.06.006