THE ROLE OF SUBTHALAMIC NUCLEUS OSCILLATORY ACTIVITY AS IT PERTAINS TO DECISION-MAKING



Baltazar A Zavala

Doctoral thesis

Nuffield Department of Clinical Neuroscience

Oxford University & St John's College

January 2015

i. Declaration

I, Baltazar Zavala, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

Signed:



Date: 22nd January 2015

ii. Abstract

The subthalamic nucleus (STN), which is the most common target for deep brain stimulation for Parkinson's disease, is known to be crucially involved in motor control. Recent appreciation of the potential non-motor side effects of STN deep brain stimulation, however, has led to speculation that the importance of this nucleus may also relate to processes involved in decision-making, particularly during high conflict scenarios. This thesis concerns itself with investigating the STN's role in action selection during conflict. I begin by recording local field potentials directly from the STN of Parkinson's disease patients while they perform a flanker task that has been shown to elicit theta (4-8 Hz) band activity in areas of the prefrontal cortex involved in cognitive control. I report that like the prefrontal cortex, the STN demonstrates elevated theta activity during conflict. I then test whether STN theta activity is related to that of the prefrontal cortex by recording from both sites simultaneously while patients perform a novel task that temporally separates conflict from stimulus onset or movement. This reveals that theta activity indeed becomes synchronized during conflict, with cortical oscillations driving those of the STN. Thirdly, I investigate how STN oscillations may affect firing rate dynamics by intra-operatively recording local field potentials and single unit activity from patients performing the flanker task. I report that both theta and beta (15-30 Hz) oscillations entrain STN neurons, but only during conflict. Finally, I record cortical and STN activity while a fourth group of patients performs the flanker task. This experiment confirms that cortico-STN theta synchrony is elevated during conflict and may also relate to across-trial adaptations to conflict and errors. Taken together these studies shed light on the mechanisms by which cortical structures may influence the STN during conflict and why STN deep brain stimulation may result in impulsivity.

iii. Contents

i. Declaration **2**
ii. Abstract **3**
iii. Contents **4**
iv. List of Figures **9**
v. List of Tables **11**
vi. List of Equations **12**
vii. Acknowledgements **13**
viii. Abbreviations **14**
ix. Publications resulting from work related to this thesis **15**
Chapter 1. Introduction **16**
1.1 Parkinson's disease, the basal ganglia, and deep brain stimulation **17**
1.1.1 Parkinson's disease 18
1.1.2 The Basal Ganglia Rate Model: The direct and indirect pathways 19
1.1.3 Deep Brain Stimulation 22
1.1.4 Beyond the rate model 24
1.1.5 Pathological oscillations in the basal ganglia 25
1.2 Computational models of the STN's role in inhibiting choices and actions **27**
1.3 Lesion and DBS studies show impaired response inhibition during conflict **29**
1.4 Electrophysiology supporting a role for the STN in response inhibition **32**
1.4.1 Firing rate studies of single neurons 32
1.4.2 Synchronized oscillations in the beta frequency (13-30 Hz) band 34
1.4.3 Synchronized oscillations in the theta frequency (4-8 Hz) band 36
1.5 Thesis objectives **37**
Chapter 2. General Materials and Methods **39**
2.1 Test Subjects and Recordings **39**
2.1.1 Ethics 39

2.1.2 Post-operative recordings in the UK 39

2.1.3 Intraoperative recordings at the NIH 40

2.2 Electrophysiology data analysis 41

2.2.1 Data pre-processing and conversion to the frequency domain 41

2.2.2 Oscillatory power analysis 43

2.2.3 Oscillatory Inter-trial phase consistency (ITPC) 48

2.2.4 Inter-site phase coherence 50

2.2.5 Inter-site Lachaux coherence 52

2.2.6 Granger Causality analysis 54

2.3 Statistical analysis 59

2.3.1 Parametric statistics 59

2.3.2 Permutations and exceedance mass testing 59

Chapter 3. Subthalamic nucleus local field potential activity during the Eriksen flanker task reveals a novel role for theta phase during conflict monitoring 62

3.1 Introduction 62

3.2 Methods 63

3.2.1 Subjects 63

3.2.2 Task 64

3.2.3 LFP data recording and analysis of power 65

3.2.4 Inter-trial phase consistency (ITPC) 66

3.2.5 Electrophysiological relationships with reaction time 67

3.3 Results 67

3.3.1 Power changes during correct trials 68

3.3.2 Theta phase inter-trial phase locking during correct trials 72

3.3.3 Power changes in error trials 77

3.3.4 Spatial focality of recordings 78

3.4 Discussion 79

3.4.1 Preliminary hypotheses for how STN spectral features help determine response timing under conditions of conflict 80

Chapter 4. Midline frontal cortex low-frequency activity drives subthalamic nucleus oscillations during conflict 84

4.1 Introduction 84

4.2 Materials and Methods 84

4.2.1 Subjects and Task 84

4.2.2 Electrophysiological data recording and analysis of power 88

4.2.3 Granger Causality analysis 89

4.2.4 Granger Causality differences between low and high conflict trials 90

4.2.5 Hilbert coherence analysis 91

4.2.6 Calculation of power and coherence distributions across STN bipolar contacts 92

4.3 Results 93

4.3.1 Behavioural results 93

4.3.2 Power changes during low and high conflict correct trials 95

4.3.3 Coherence differences between EEG and STN in low and high conflict trials 97

4.4 Discussion 101

4.4.1 Role of theta-delta activity in conflict 103

Chapter 5. Human subthalamic nucleus theta and beta oscillations entrain neuronal firing during sensorimotor conflict 106

5.1 Introduction 106

5.2 Methods 107

5.2.1 Intraoperative recordings during deep brain stimulation surgery 107

5.2.2 Flanker task 108

5.2.3 Electrophysiology 109

5.2.4 LFP power and phase analysis 110

5.2.5 Spiking activity analysis 111

5.2.6 Spike-phase interactions 113

5.3 Results 116

5.3.1 Task and Behavior 116

5.3.2 Macroelectrode LFP power and phase changes 117

5.3.3 Microelectrode single unit firing changes 120

5.3.4 LFP-single unit interactions during conflict 126

5.4 Discussion 131

Chapter 6. Human subthalamic nucleus-medial frontal theta phase coherence is involved in conflict and error related cortical monitoring 136

6.1 Introduction 136

6.2 Methods 137

6.2.1 Subjects and task 137

6.2.2 LFP data recording and analysis of power 139

6.2.3 Intertrial phase consistency (ITPC) 142

6.2.4 Intersite phase coherence 143

6.3 Results 144

6.3.1 Behavioural Effects 144

6.3.2 Conflict related difference in STN LFP and mesial frontal theta activity 145

6.3.3 Post Error related difference in STN LFP and mesial frontal EEG theta phase coherence 149

6.3.4 Pre-stimulus correlates of conflict behaviour 152

6.4 Discussion 153

6.4.1 Conflict related STN theta power and phase play separate roles within a trial 154

6.4.2 STN theta activity also plays a role across trials 155

Chapter 7. Discussion 158

7.1 Potential limitations of this thesis 158

7.2 A hypothesized role of STN theta band power during action selection..... 161

7.3 A hypothesized role of STN theta phase during action selection, conflict, and post error monitoring 163

7.4 The STN is a point of convergence in a larger response inhibition network 164

7.5 Paradoxes of deep brain stimulation of the STN 167

7.6 Future Perspectives 171

Chapter 8. References..... 174

iv. List of Figures

Figure 1.1. Basal ganglia and the rate model. 20

Figure 1.2. Human electrophysiology supporting a role for the STN in decision-making. 34

Figure 2.1. Electrophysiology data processing for power analysis. 46

Figure 2.2. Inter-trial phase consistency (ITPC). 49

Figure 2.3. Granger connectivity analysis. 58

Figure 3.1. Flanker task. 68

Figure 3.2. Effects of conflict on LFP. 69

Figure 3.3. Differences between slow and fast trials following high conflict cues. 71

Figure 3.4. Cue-locked theta phase realignment is disrupted in slow-high conflict trials. 74

Figure 3.5. Cue-locked phase alignment correlates with reaction time. 77

Figure 3.6. LFP in error trials across all subjects. 78

Figure 4.1. Task and behaviour. 94

Figure 4.2. Power changes in STN LFP. 96

Figure 4.3. Changes in cortico-STN LFP coherence. 98

Figure 4.4. Changes in continuous cortico:STN 2-8 Hz coherence using Hilbert transform. 99

Figure 4.5. Changes in directed coherence between Fz-Cz and STN LFP. 100

Figure 4.6. Finer frequency resolution directed coherence between Fz-Cz and STN LFP. 101

Figure 5.1. Flanker task during intraoperative STN recordings. 116

Figure 5.2. STN theta band activity is higher during conflict. 118

Figure 5.3. Response-aligned STN LFP activity. 120

Figure 5.4. Task related increased spiking activity in STN cells. 122

Figure 5.5. Conflict related differences in increased spiking activity in STN cells. 123

Figure 5.6. Conflict related differences in rate of increase in spiking activity in STN cells 125

Figure 5.7. Conflict related differences in downward firing STN cells. 126

Figure 5.8. Theta and beta oscillations entrain STN cells during conflict. 127

Figure 5.9. Time evolving theta and beta spike phase locking in STN cells. 129

Figure 5.10. Early and late cell differences in spike phase locking. 130

Figure 6.1. Task, behaviour, and theta. 145

Figure 6.2. Conflict related theta activity in the STN and mPFC. 147

Figure 6.3. Theta STN-mPFC coherence is higher on trials that follow an error. 151

Figure 6.4. Inter-trial STN theta power is related to slower reaction times during conflict..... 153

v. List of Tables

Table 3.1. Clinical details from first Flanker experiment 63

Table 4.1. Clinical details from dots experiment. 85

Table 5.1. Clinical details from second Flanker experiment. 107

Table 6.1. Clinical details from third Flanker experiment 138

vi. List of Equations

Equation 2.1. Inter trial phase consistency (ITPC).....48

Equation 2.2. Phase coherence 50

Equation 2.3. Traditional coherence.....52

Equation 2.4. Lachaux coherence.....53

vii. Acknowledgements

I would like to thank first and foremost, my two phenomenal supervisors, Peter Brown and Kareem Zaghloul. They took me into their labs when all I knew about the brain was that it contained neurons and all I knew about Matlab was how to plot a figure. I could not have asked for better guidance when going through the research process and better support when my preliminary experiments (TMS & TACS) all failed. The example they set in both the clinic and the laboratory has been truly inspiring, and I hope to one day be half as extraordinary as they are.

I would also like to thank all of the members of the Brown and Zaghloul labs for their invaluable friendship as well as their tremendous generosity. Without their endless contributions of time, experience, support, tacos, pints, and hand-me-down baby clothes, I can confidently say I would not have made it through my DPhil.

The opportunities to conduct such interesting research would also not have been possible without the many patients who have volunteered to let me record the activity deep inside their brain. Their energy, optimism, and desire to fight back in any way they can (even if it involves doing the dots task) has continually fueled my interest to pursue a career in medical research.

I am also grateful to the Rhodes trust, the Oxford-Cambridge Scholars program, and the Medical Scientist Training Program for giving me the opportunities to pursue my passions and learn so much. Finally, I would like to thank my wife and daughter who's loving smiles have reminded me every single day (particularly those involving negative results and harsh reviewers) how truly blessed I am.

This thesis is dedicated to the man who made me who I am, R. John Arrambide (26th September 1961 - 20th January 2015).

viii. Abbreviations

ACC	Anterior cingulate cortex
ANOVA	Analysis of variance
DBS	Deep brain stimulation
DLPFC	Dorsolateral prefrontal cortex
EEG	Electroencephalogram
GABA	Gamma-Aminobutyric acid
GPi	Globus pallidus interna
GPe	Globus pallidus externa
IFC	Inferior frontal cortex
ITPC	Inter-trial phase consistence
LFP	Local field potential
mPFC	Medial prefrontal cortex
MPTP	1-methyl-4-phenyl-1,2,3,6-tetrahydropyridine
NIH	National Institutes of Health
PD	Parkinson's disease
Pre-SMA	Pre-supplementary motor area
SEM	Standard error of the mean
SNc	Substantia nigra pars compacta
STN	Subthalamic nucleus
TACS	Transcranial alternating current stimulation
TMS	Transcranial magnetic stimulation
UPRDS	Unified Parkinson's Disease Ratings Scale

ix. Publications resulting from work related to this thesis

1. Zavala B, Brittain J-S, Jenkinson N, Ashkan K, Foltynie T, Limousin P, Zrinzo L, Green AL, Aziz T, Zaghloul K, Brown P. (2013) Subthalamic Nucleus Local Field Potential Activity during the Eriksen Flanker Task Reveals a Novel Role for Theta Phase during Conflict Monitoring. J Neurosci 33:14758–14766.

2. Zavala BA, Tan H, Little S, Ashkan K, Hariz M, Foltynie T, Zrinzo L, Zaghloul KA, Brown P. (2014) Midline Frontal Cortex Low-Frequency Activity Drives Subthalamic Nucleus Oscillations during Conflict. J Neurosci 34:7322–7333.

3. Zavala B, Zaghloul KA, Brown P. (2015) The Subthalamic Nucleus, Oscillations, and Conflict. Mov Disord. 30:328-338

4. Zavala B, Damera S, Dong J, Lungu C, Brown P, Zaghloul KA. Human Subthalamic Nucleus Theta and Beta Oscillations Entrain Neuronal Firing During Sensorimotor Conflict. Under review.

Chapter 1. Introduction

Being able to make decisions both quickly and accurately is a key skill. In the wild, a well-balanced speed-accuracy trade off in decision-making greatly increases the chances of survival. In daily human life, these same mechanisms can often make the difference between success and failure. One fundamental requirement of successfully balancing speed-accuracy trade-offs is the ability to alter the amount of time dedicated to a decision based on the difficulty of that decision. If your boss asks you if you would like an increased salary, you may want to respond quickly before they change their mind. If instead you are asked if you would like a 5% raise now or a 10% raise in two years, you may want to think for a little bit longer before taking the immediate increase. As this example illustrates, responding impulsively during conflict can result in less than ideal outcomes that can often have a negative impact on our lives. Indeed many pathological disorders such as attention deficit hyperactivity disorder, obsessive compulsive disorder, and even schizophrenia are associated with impulsivity (Lipszyc and Schachar, 2010).

Over the past two decades, there has been a growing clinical and scientific interest in the mechanisms by which the brain adjusts the amount of time taken to make a decision as well as the disruptions of the brain that lead to impulsive decision-making. Often times, these mechanisms are probed by simple response inhibition tasks that ask subjects to either rapidly inhibit a preplanned action or to select an action based on conflicting sensory information. Through these studies, large progress has been made in identifying several cortical regions that are involved in preventing impulsivity (Botvinick et al., 2004). Recently, however, computational models have increasingly implicated the basal ganglia as the final point of convergence for conflict triggered response inhibition. Specifically, the models have led to the hypothesis that the subthalamic nucleus (STN), which is the most common target for deep brain stimulation (DBS) for Parkinson's disease (PD), is a key node in the brain network responsible for "holding your horses" during conflict (Frank, 2006; Frank et al., 2007). In line with this hypothesis, several

studies have shown that STN DBS impairs response inhibition (see section 1.3 below), cross-sectional studies have shown that STN DBS is associated with higher scores on impulsivity rating questionnaires (Hälbig et al., 2009) and anecdotal reports have suggested that on rare occasions STN DBS may even trigger pathologically impulsive behaviour such as hyper-sexuality and pathological gambling (Romito et al., 2002; Herzog et al., 2003; Lu C et al., 2006; Mandat et al., 2006; Smeding et al., 2007). Understanding the mechanisms underlying the STN's role in cognition, particularly the interactions to be had between the STN and diverse cognitive areas of cortex, is critical if we are to avoid cognitive side-effects in treating motor disorders, or, indeed, turn them to our benefit in the treatment of psychiatric diseases. In this thesis, I will first review the literature that has previously implicated the STN in conflict triggered response inhibition. I will then present the results from several of the studies we conducted in which we found that theta band oscillations (4-8 Hz) synchronize the activity of the STN to areas of prefrontal cortex that are thought to detect conflict. Finally, I will discuss how the disruption of these networks by DBS may result in impulsivity and what implications these findings may have for future research and therapies.

1.1 Parkinson's disease, the basal ganglia, and deep brain stimulation

As the STN is primarily known for the role it plays as the primary target for therapeutic DBS for PD, it is important to first review several of the fundamental breakthroughs that have been made by studying the STN in the context of PD. Efforts to understand PD have not only shed invaluable light on the mechanisms that lead to several pathological impairments; they have also provided incredible insight into the roles the basal ganglia play in motor control, learning, and decision-making. As these insights set the foundation on which many of the computational models of STN and decision-making are based, I will begin by briefly reviewing much of what is known about the basal ganglia and PD.

1.1.1 Parkinson's disease

Parkinson's disease (PD) is currently the second most common neurodegenerative disease, affecting approximately 7 million individuals world-wide (prevalence= 0.3% in entire population and 1% in individuals over 60 years of age) (de Lau and Breteler, 2006; Yao et al., 2013). It is primarily characterized by an abnormal slowness of movements (bradykinesia), rigidity, and resting state tremors in the 4-6 Hz oscillatory range (Hughes et al., 1992). Other motor symptoms of PD can also include temporary inabilities to initiate movements (freezing), postural instability, and impaired speech and swallowing. As the disease progresses with time, the severity of the symptoms increases and often extends to non-motor impairments including memory loss, depression, and sleep disorders.

Though the causes of PD remain highly debated and likely reflect a combination of genetic and environmental factors (Tanner CM et al., 1999; Lesage and Brice, 2009), the disease is primarily characterized by a loss of mid-brain dopamine producing neurons in the substantia nigra pars compacta (SNc) (Hassler, 1938). Accordingly, the strategy most commonly used to treat PD is to attempt to restore the levels of dopamine to their pre-symptomatic levels. This is predominantly achieved by administering to patients the dopamine precursor Levodopa, which is subsequently converted to dopamine once inside the brain (Birkmayer and Hornykiewicz, 1961; Hornykiewicz, 2010). Levodopa is often times also supplemented with several other medications such as dopamine receptor agonists like Apomorphine, Ropinirole, and Pramipexole or dopamine breakdown inhibitors like Rasagiline. Though these treatments initially produce very significant improvements in symptoms, over time many patients become less responsive to dopaminergic therapy. Many also experience an increase in fluctuations in their level of responsiveness to the drug, leading to unpredictable On and Off periods of their symptoms. Finally, Levodopa and dopaminergic agonists can often induce undesirable side affects such as rapid, uncontrollable movements (dyskinesias) or even impulse control disorders (Voon V and Fox SH, 2007;

Rodriguez-Oroz et al., 2011), which will be discussed further below (see section 7.5 “Paradoxes of deep brain stimulation of the STN”). Once these symptoms become too severe, patients become strong candidates for more invasive therapies like DBS, which also will be discussed further below (section 1.1.3 “Deep Brain Stimulation”).

Though the relationship between SNc neuronal loss and PD is well established, how these microscopic changes in dopamine levels result in the symptoms associated with PD remains poorly understood. Intriguingly, several studies have shown that dopaminergic loss at the time of diagnosis can reach up to levels of 70 % (average estimates of dopaminergic loss at symptom onset range from 50-70%) before any symptoms become apparent (Marsden, 1990; Damier et al., 1999; Rajput et al., 2008; Kordower et al., 2013). This raises the possibility of compensatory mechanisms that become saturated as dopaminergic loss increases. Moreover, it is likely that the loss of SNc neurons per se may not be what directly leads to pathological symptoms. Rather the gradual, network-wide changes in connectivity and reactivity that result from chronic dopaminergic depletion may be responsible. This idea has been strongly supported by many of the findings made from studying primates that have been rendered Parkinsonian by dopaminergic depletion via 1-methyl-4-phenyl-1,2,3,6-tetrahydropyridine (MPTP) (Burns et al., 1983; Langston et al., 1983). One of the most influential of these was the finding that chronic dopamine depletion results in an increase in baseline firing rates of the output structure of the basal ganglia, the globus pallidus internus (GPi), as well as its up stream target, the STN (Miller and DeLong, 1987; Filion and Tremblay, 1991). This finding, combined with studies of the structural anatomy of the basal ganglia resulted in what is now commonly referred to as the “rate model” of basal ganglia architecture and to a firing rate based explanation of PD (Albin et al., 1989; DeLong, 1990).

1.1.2 The Basal Ganglia Rate Model: The direct and indirect pathways

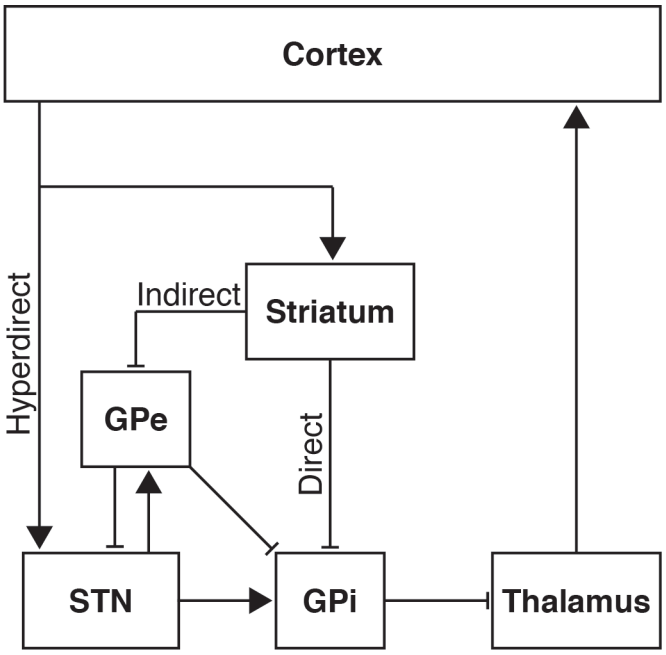


Figure 1.1. Basal ganglia and the rate model.
Pre-motor cortical inputs project to the striatum where they can either activate a movement by inhibiting the GPi via the direct pathway or inhibit a movement by inhibiting the GPe via the indirect pathway. Cortical inputs also project to the STN via the fast conducting hyper-direct pathway.

The traditional rate modal of basal ganglia architecture posits that there are two distinct pathways by which the basal ganglia can affect movement (Albin et al., 1989; DeLong, 1990). Both pathways can be activated by an increase in cortical inputs to the striatum, which sends GABAergic projections to the internal and external segments of the globus pallidus (GPi and GPe, respectively, Figure 1.1). According to the model, if the cortical inputs activate an area of the striatum that projects to the GPi, a pro-kinetic "Go" signal is thought to be activated as increased output from the striatum will result in an inhibition of the GPi, which will then result in a release of inhibition of the thalamus. Activation of the areas of the striatum that project to the GPe, on the other hand, will result in an anti-kinetic "NoGo" signal as the GPe inhibits the GPi both directly via GABAergic projections and indirectly via inhibition of the STN. Therefore, an inhibition of the GPe by striatal NoGo inputs will lead to a net increase in both STN and GPi activity and thus increased inhibition of the thalamus. Thus when an individual wishes to make a movement, the model posits that cortical activity related to the movement activates the "direct "

Go pathway to trigger the movement. Conversely, when an individual wishes to inhibit a movement or other movement alternatives (thereby assisting movement selection), cortical activity activates the "indirect" NoGo pathway.

In the healthy basal ganglia, dopamine is thought to play a central role in controlling motor learning by altering the balance between the Go and the NoGo pathways. The cortico-striatal synapses of GPi projecting striatal neurons ("Go" neurons) predominantly express D1 dopamine receptors (Gerfen, 1992; Ince et al., 1997). These receptors can result in long term potentiation of a synapse when activated (Kitai et al., 1976; Hernández-López et al., 1997). In contrast, the cortico-striatal synapses of the GPe projecting striatal neurons ("NoGo" neurons) predominantly express D2 dopamine receptors (Gerfen, 1992; Ince et al., 1997), which can result in long term depression of a synapse when activated (Hernández-López et al., 2000). Thus, task related dopaminergic increases in the basal ganglia result in a strengthening of the Go pathway and a weakening of the NoGo pathway. This is thought to contribute towards motor learning because dopaminergic SNc cells that project to the striatum demonstrate an increase in firing when subjects are rewarded (Schultz et al., 1997; Schultz, 2002; Zaghloul et al., 2009). This increase is thought to enhance the Go pathways that triggered the rewarded action and suppress the NoGo pathways that would have interfered with its execution (Frank et al., 2004; Frank, 2005). In contrast, when an individual is punished, tonic dopamine levels drop (Schultz, 2002; Satoh et al., 2003), resulting in a suppression of the Go pathway associated with the incorrect action and an enhancement of the NoGo pathway that would have prevented its execution.

As dopamine is thought to be crucial to the maintenance of a properly balanced Go-NoGo circuitry, the rate model predicts that the chronic loss of dopamine experienced in PD results in an imbalance of the "Go" and "NoGo" pathways as the "Go" pathway can no longer be strengthened and the "NoGo" pathway can no longer be suppressed (Albin et al., 1989; DeLong, 1990). In line with these predictions, the strengthening of the NoGo pathway is thought to

underlie many of the akinetic symptoms experienced by PD patients such as bradykinesia, rigidity, and freezing. Further support for the rate model's predictions is the aforementioned finding that dopaminergic denervation in primates leads to an increase in tonic GPi output (consistent with a hypo-active Go pathway) as well as a decrease in tonic GPe output (consistent with a hyper-active NoGo pathway) (Miller and DeLong, 1987; Filion and Tremblay, 1991). This results in a chronic inhibition of the thalamus, which is further exacerbated by the effects dopamine depletion has on the STN. As the STN receives feedback inhibition from the GPe, the strengthening of the NoGo pathway that results in decreased GPe firing is thought to also lead to increased firing of the STN. As the STN projects to the GPi, this is thought to lead to even higher GPi firing levels and thus a further inhibition of movement.

1.1.3 Deep Brain Stimulation

Soon after the discovery that Parkinsonian primates demonstrated elevated firing in the GPi and the STN, Delong and colleagues made the landmark discovery that a lesion to the STN could reduce the firing of the GPi (Wichmann et al., 1994) and drastically improve the symptoms experienced by the Parkinsonian primate (Bergman et al., 1990). This work further validated the rate model hypothesis of PD and ushered in a new era of research in PD and the basal ganglia. Though the finding that GPi lesions could result in improvements in PD symptoms had been known as early as the 1950s (Narabayashi and Okuma, 1953; Svennilson et al., 1960; Laitinen et al., 1992), the high risks associated with deep brain lesions, together with the discovery of Levodopa as a successful therapy, had temporarily overshadowed any further investigations in surgical treatments for PD (Wichmann et al., 2011). The formalization of the rate model, however, combined with an increased awareness of the long term drawbacks of Levodopa, resulted in a new interest in identifying ways of altering the abnormal basal ganglia activity observed in PD. Initially this was accomplished by a reintroduction of pallidotomies in which the GPi was thermally lesioned by a strong injection of current (Laitinen et al., 1992). This operation

was very similar to an operation that was being used at the same time to lesion the thalamus as a treatment for essential tremor. Though the thalamic lesion had also been used for quite some time, it suffered the same drawbacks of STN and GPi lesions. These included risks for haemorrhage during the surgery, difficulty in controlling the size of the lesion, and irreversible side effects that could result from the lesion. These drawbacks encouraged Benabid and colleagues to search for reversible ways of altering the output of the surgical targets, and in 1987 they made the seminal discovery that chronic high frequency stimulation of the thalamus resulted in a very similar clinical effect as the thalamic lesion (Benabid et al., 1987, 1991). Importantly, however, DBS had the added benefits of being less invasive and being reversible. Soon after Benabid et al discovered the efficacy of thalamic DBS for essential tremor, DeLong et al's discovery that STN lesions could improve Parkinsonian symptoms motivated others to try DBS of the STN for PD. This procedure, first done in MPTP primates in 1993 (Benazzouz et al., 1993) and in humans in 1995 (Limousin et al., 1995), led to what is now considered to be the golden standard of functional neurosurgery. Not only can STN DBS improve many of the symptoms of PD similar to a lesion, it can also alleviate many of the side effects of dopaminergic medication like dyskinesias and motor fluctuations. Importantly, it does so in a reversible manner and many negative side effects that result from DBS can often be eliminated by reducing the amount of current delivered or altering the stimulation parameters used.

The discovery of DBS ushered in a new era of treatments for patients suffering from PD. Typically, DBS surgery involves two procedures (Foltynie and Hariz, 2010). In the first, an electrical wire is inserted inside the basal ganglia and small microelectrodes at the tip of the wire are used to record action potential firing of neurons within the basal ganglia. These wires have been used to explore many of the firing rate dynamics of the basal ganglia as well as how these dynamics are altered during movement, learning, and many other actions (Magariños-Ascone et al., 2000; Rodriguez-Oroz et al., 2001; Williams et al., 2005; Zaghloul et al., 2009, 2012; Hanson

et al., 2012). After the STN has been localized by its distinct firing pattern, the microelectrode is replaced with a larger macroelectrode that is capable of stimulating the STN as well as recording Local Field Potential (LFP) activity of the STN. These LFP recordings predominantly measure the fluctuations in synaptic inputs to the recorded area (Buzsáki et al., 2012). Following the implantation of the macroelectrode in the STN, a second surgical procedure is done in which the DBS electrode is connected to a subcutaneous battery pack that is inserted in the subclavicular fossa. This second procedure can take place during the same operation in which the electrode is implanted in the STN or in a second operation up to a week later. When the battery implantation is delayed, the patients remain in the hospital with their DBS leads externalized. During this time, physicians and researchers are able to work with the patients to record the LFP activity of the STN as the patients perform various tasks.

1.1.4 Beyond the rate model

The externalized electrode recordings have allowed researches a rare window into the brain that has been greatly utilized over the past two decades. Initially, the clinical outcome similarities between the STN lesion and STN DBS led to the hypothesis that DBS acts as a virtual lesion that shuts down the STN firing activity and thus restores basal ganglia neuronal outputs to less disruptive levels (Benazzouz et al., 1995; Limousin et al., 1998). This interpretation was very much in line with the rate model prediction that the increases in GPi firing that followed dopaminergic lesions were responsible for the anti-kinetic symptoms of PD. Over the past few years, however, several lines of evidence have seriously questioned the validity of this interpretation and of the rate model in general (Marsden and Obeso, 1994; Brown and Eusebio, 2008). First, it has been shown using several modalities that DBS does not lesion or inhibit the STN, but rather it actually activates the STN output and results in an increase in GPi firing (Windels et al., 2000; Hashimoto et al., 2003; Stefani et al., 2005; McIntyre and Hahn, 2010). This finding not only goes against the lesion hypothesis of DBS function, it also seems to

contradict the rate model entirely by suggesting that increased GPi firing is associated with an improvement of motor symptoms. Other shortcomings of the rate model have been further revealed by studies involving lesions to the GPi or thalamus. The rate model would predict that lesions to the GPi would result in a hyper-kinetic activity such as dyskinesia and that lesions to the thalamus would result in akinesia. Neither of these predictions have been true (Marsden and Obeso, 1994; Obeso et al., 1997; Brown and Eusebio, 2008), leading some to speculate that other factors (see next section 1.1.5 "Pathological oscillations in the basal ganglia" below) beyond just the firing rate changes are involved in PD (Brown, 2003, 2007; Hammond et al., 2007). Furthermore, it has been repeatedly shown that prior to movement, individual neurons of the GPi actually show an increase in firing (reviewed in Goldberg et al., 2013). As the rate model would predict a decrease in firing prior to movement, these findings suggest that even simple motor control may be more complicated than just a Go and a NoGo pathway (Nambu et al., 2002; Goldberg et al., 2013).

Despite the rate model's shortcomings, it important to note the many great theoretical and clinical advances that resulted from the predictions it made. Furthermore, optogenetic studies have shown that the direct and indirect pathways do activate and inhibit movement, respectively (Kravitz et al., 2010) and that dopaminergic activation of the two pathways does lead to long term potentiation and depression, respectively (Kitai et al., 1976; Hernández-López et al., 1997, 2000). In the words of one researcher describing the evidence against the rate model, "While these data certainly challenge the simplest version of the direct/indirect model, one need not 'throw the baby out with the bath water' and reject the standard model altogether, especially given its contributions to many theoretical and practical advances" (Frank, 2006).

1.1.5 Pathological oscillations in the basal ganglia

Soon after the discovery that dopaminergic lesions in the MPTP primate resulted in elevated firing rates in the STN and GPi, several studies showed that the patterned activity of the basal

ganglia was also changed. Whereas neuronal activity in healthy primates fires tonically, PD primates show a shift towards burst firing in the GPe, STN, and GPi (Miller and DeLong, 1987; Filion and Tremblay, 1991; Bergman et al., 1994; Wichmann and Soares, 2006). This effect has been subsequently shown to also be true in the human basal ganglia (Levy et al., 2000). Furthermore, the activity of the entire motor network seems to become more synchronized in an oscillatory manner following the dopaminergic lesion. Using recordings made directly from the human basal ganglia, Brown and colleagues have repeatedly established an association between exaggerated oscillatory "beta" band activity (roughly 10-30 Hz) and Parkinson's disease (Brown, 2003, 2007; Hammond et al., 2007). Spectral activity from the resting state Parkinsonian brain demonstrates a peak in the beta band activity (Brown et al., 2001). This peak is reduced when patients are given their medication and the degree of beta suppression in the STN and in the cortex correlates with the improvement of motor symptoms (Silberstein et al., 2005; Kühn et al., 2006; Ray et al., 2008). Similarly, DBS has been shown to suppress beta band activity, and this too correlates with motor improvements (Brown et al., 2004; Silberstein et al., 2005; Eusebio et al., 2011). Off medication, beta band activity seems to entrain individual STN neuronal firing (Kühn et al., 2005; Weinberger et al., 2006) and to tightly couple the STN to the GPi (Brown et al., 2001).

The work related to the beta band and PD has lead to the hypothesis that akinetic PD symptoms are not simply due to the increased firing of the basal ganglia structures but rather they are caused by an excessive synchronization at a frequency that is anti-kinetic (Brown, 2003, 2007; Hammond et al., 2007). The relationship between excessive beta synchrony and akinetic PD symptoms supports this hypothesis, as do the findings that PD patients and healthy human subjects show a decrease in beta band activity in the motor network during movement. Movement related beta band decreases have been shown at the single-unit (Levy et al., 2002; Amirnovin et al., 2004), LFP (Cassidy et al., 2002; Kühn et al., 2006), and cortical levels (Alegre et al., 2003),

and stimulation of the STN at the beta band has been shown to slow movement (Chen et al., 2007, 2011; Eusebio et al., 2008). This has lead to the hypothesis that beta may reflect a "status quo" signal (Engel and Fries, 2010). In light of this interpretation, Brown et al have proposed that DBS may function by disrupting the excessive beta synchrony and thus allowing the basal ganglia and the rest of the motor networks to encode movement related activity (Hammond et al., 2007).

Unfortunately, current stimulation protocols are not able to selectively suppress only the beta band activity as they involve very high frequency stimulation that is thought to disrupt all patterned output of the STN (Eusebio et al., 2011; Moran et al., 2012; Whitmer et al., 2012). This disruption includes not only the pathological beta band activity, but also the normal physiological activity that might play an important role in various motor and cognitive processes. As the roles of the STN oscillatory frequencies outside of the beta and gamma ranges remain relatively understudied, the relationship between DBS, oscillations, and unwanted symptoms of DBS remains unclear.

1.2 Computational models of the STN's role in inhibiting choices and actions

Over the past few years, some of the side effects that are associated with STN lesions and DBS have motivated new lines of research into what role the STN normally plays in the healthy basal ganglia. Though the STN was initially thought to be no more than a complementary part of the indirect NoGo pathway, it is now well known that the STN serves as a crucial secondary point of entry for cortical input to the basal ganglia (Monakow et al., 1978; Nambu et al., 2000; Tokuno and Nambu, 2000). This "hyperdirect" pathway that connects the cortex directly to the GPi via the STN is thought to conduct much faster than the cortical-striatal inputs of the direct pathway (Nambu et al., 2002). Until recently, however, the functional role of the hyper-direct pathway was relatively unexplored.

As mentioned above, the primary down-stream target of excitation from the STN is the GPi, which sends inhibitory projections to the motor thalamus (Figure 1.1). Thus, in the context of the traditional rate model of basal ganglia connectivity (Albin et al., 1989; DeLong, 1990), increased firing in the STN is expected to increase the activity of the GPi, which would then lead to disfacilitation of the cortex by the thalamus and thus an inhibition of movement. This view, though widely acknowledged as an over simplification of the basal ganglia (see section 1.1.4 "Beyond the rate model" above), has inspired two major computational models of STN function during decision-making (Frank, 2006; Bogacz and Gurney, 2007). Both models posit a crucial role for the STN in response inhibition under conditions in which several responses are entertained, as when conflicting information is presented to the subject. Though the mathematical relationships underling the dynamics of STN activity differ between the models, they both predict that the hyperdirect cortico-STN pathway is responsible for inhibiting responses during conflict to prevent impulsive actions. They both also, at least by inference, extend this function to inhibiting the wrong choice during conflict, thus implicating the STN in decision-making.

The hypothesized dynamics underlying the models can be summarized as follows. During straightforward action selection that does not involve conflict, the pre-motor cortical activity associated with the correct response initially activates the STN via the fast conducting hyper-direct pathway (Figure 1.1). This activation results in a temporary inhibition of movements due to the broad excitation of GPi by STN neurons (Hazrati and Parent, 1992; Nambu et al., 2002). After a brief delay, however, the feedback inhibition from the GPe suppresses the STN's net anti-kinetic effect, and cortico-striatal activity corresponding to the correct response inhibits GPi via the "direct" pathway. Once the cortico-striatal activity reaches a certain threshold, the tonically active GPi neurons are inhibited sufficiently to release thalamic facilitation of the cortex.

During situations involving conflict, on the other hand, the models assume that two or more pre-motor responses are simultaneously activated by the conflicting options. This leads to more

marked activation of the STN, and thereby GPi (Figure 1.1). In order to overcome this conflict triggered increase in global inhibition, any movement-related cortico-striatal input will need to be stronger to silence the GPi. This effectively results in a higher "threshold" of cortical evidence needed to drive a response. Over time, the pre-motor inputs associated with the correct response will continue to rise, while those associated with the incorrect response will not, eventually resulting in the selection of the correct response. This process will take time, however, resulting in longer reaction times during high conflict trials relative to low conflict trials. Despite the longer reaction times, this proposed process would allow for improved action selection by preventing the selection of the incorrect response. Without the conflict-triggered increase in global inhibition, the incorrect response would occasionally be selected due to either noise in the system or to pre-potent biases towards the incorrect response. As the STN is crucial in inhibiting responses when two or more potential responses are simultaneously entertained, these models assign to the STN the role of "holding your horses" during high conflict decisions. This has been formalized in the context of the drift diffusion model of decision-making. In this model, a noisy neuronal process accumulates information supporting one of two alternatives until the evidence difference between the two alternatives is large enough to cross a "decision threshold" and execute the winning response (Ratcliff and McKoon, 2008). During conflict, inhibitory activity from the STN adjusts the evidence threshold the winning alternative will need in order to win the drift diffusion race (Cavanagh et al., 2011; Ratcliff and Frank, 2012; Green et al., 2013; Obeso et al., 2014). According to these predictions, the disruption of STN activity via DBS should result in impulsive behaviour in the face of conflict.

1.3 Lesion and DBS studies show impaired response inhibition during conflict

Prior to the predictions of the computational models described above, there were several lines of evidence already supporting the claim that the STN is crucial for response inhibition. In a delayed

reaction time task where rats were rewarded if they withheld responses until a cue was presented, but not if they responded too early or too late, dopaminergic lesions led to a delay in the response times, similar to that seen in PD patients and resulting in more trials that were missed due to late responses (Baunez et al., 1995). A subsequent lesion of the STN led to a speeding in reaction times, consistent with the effects of DBS on PD. Following the STN lesion, however, the rodents were impaired in their ability to withhold the response until the cue was presented. Thus although the STN lesion produced DBS-like improvements in motor control, it came at the cost of impulsivity in this simple task.

In line with lesion studies in rodents and the predictions made by the computational models, PD patients receiving DBS of the STN have generally been shown to respond more impulsively in the laboratory setting (Jahanshahi et al., 2000; Witt K et al., 2004; Frank et al., 2007; Ballanger et al., 2009; Hershey et al., 2010; Cavanagh et al., 2011; Coulthard et al., 2012; Green et al., 2013; Antoniades et al., 2014) (but also see van den Wildenberg et al., 2006; Swann et al., 2011; Mirabella et al., 2012 and section 7.5 "Paradoxes of deep brain stimulation of the STN"). Most of these studies have focused on scenarios in which a fast, pre-potent response must be suppressed in favour of a slower, more controlled response. Though these tasks usually do not require an overt decision, conflict arises between the activity related to the primed incorrect response and that related to the unprimed correct response. Consider, for example, performance on the "Stroop" task. Here subjects must specify the colour of the font of printed words, whilst suppressing the pre-potent drive to read the word, which may be for a different colour (i.e. the word "red" written in blue font requires the response "blue" (Stroop, 1935)). PD patients with STN electrodes have more errors when their stimulators are turned on (Jahanshahi et al., 2000; Witt K et al., 2004). Likewise, when patients perform a Go-NoGo task where they receive a go stimulus on most trials, priming them to respond quickly, but are occasionally presented with a NoGo stimulus in which they must cancel the planned Go response, subjects commit more errors

in the DBS ON condition (Ballanger et al., 2009; Hershey et al., 2010). In line with these results, several other studies using simple action selection tasks involving conflict have shown that STN DBS as well as subthalamotomies induce more errors by impairing response inhibition (Ray et al., 2009; Green et al., 2013; Antoniades et al., 2014; Obeso et al., 2014).

Several studies have also suggested that the STN may also be involved in more complicated, overt decisions (Frank et al., 2007; Cavanagh et al., 2011; Coulthard et al., 2012). As a follow up to their modelling work, Frank et. al demonstrated that when DBS stimulators were turned off, patients behaved no differently than healthy controls in a probabilistic paired choice decision task (Frank et al., 2007). This task involves a choice between two symbols that are associated with differing probabilities of reward. On low conflict decisions, the difference between the two options is sufficiently large that the correct choice is obvious (80% vs. 20%). On difficult decisions, the probability difference is close enough to induce conflict (70% vs. 60%). As would be expected, healthy controls and PD patients in the off DBS state responded significantly more slowly during the high conflict decisions relative to the low conflict decisions. When the STN stimulators were turned on, the patients responded faster during both low and high conflict trials, consistent with the motor improvements associated with DBS. However, in the DBS on condition, the patients actually responded significantly faster during the high conflict trials than they did during the low conflict trials, particularly during win-win scenarios. These differences could not be explained by the motor improvements alone, as faster movements on DBS would not explain why the conflict induced slowing would become conflict induced speeding with DBS. As expected, the faster responses led to suboptimal decision-making and more errors. Similarly, when patients are asked to integrate several lines of evidence to make a decision, STN DBS impairs the usual slowing of reaction time that occurs when the final piece of evidence they receive is contradictory to the previously received evidence (Coulthard et al., 2012). Taken together, the rodent lesion and human DBS studies show that the STN plays an important role in

response inhibition, which can under some circumstances influence decision-making. The principle focus of this thesis is to study the mechanisms by which the STN might carry out this function. Prior to our work however, several electrophysiology studies had begun to study how the STN might be involved in response inhibition.

1.4 Electrophysiology supporting a role for the STN in response inhibition

1.4.1 Firing rate studies of single neurons

The majority of single neuron studies involve animal models, and so have mostly been restricted to simple response inhibition paradigms. In one study, when non-human primates were primed to make saccades in a given direction and then asked to inhibit that response in favour of making a saccade in the opposite direction, STN neurons exhibited increased firing activity on the switch trials (Isoda and Hikosaka, 2008). More recently, others have attempted to tie the increased firing activity seen in the STN during response inhibition to the activity observed in the rest of the basal ganglia by simultaneously recording neuronal activity in the rodent striatum, GPi, GPe, and STN during a stop signal reaction time task (Schmidt et al., 2013). This task involves presenting a go stimulus that is occasionally followed by a stop stimulus requiring the cancellation of the initiated movement (Logan et al., 1984). Schmidt et al. found that during the stop trials, neurons of the STN always increased their firing after the stop cue. Only on successful stop trials, however, did neurons in the down-stream GPi increase their firing rate. On the failed stop trials, the inhibitory activity of the cortico-striatal direct pathway inputs managed to suppress the firing rate of GPi neurons leading to the execution of the response (see Figure 1.1). This work is in line with the predictions made by the rate model and provided compelling evidence that response inhibition involves a race between the response-driving activity of the striatum and the braking activity of the STN.

In order to show that STN activity also plays a role in overt decision-making in humans, Zaghloul et al. recorded from the STN of PD patients while they were undergoing DBS implantation surgery (Zaghloul et al., 2012). These recordings were made while subjects performed a probabilistic paired choice task similar to the one used by Frank et al (Frank et al., 2007). In line with the animal studies, STN neurons showed increased firing during the high conflict trials in which the two options contained a similar probability of reward (Figure 1.2A). Furthermore, the firing rate of the STN neurons correlated with the amount of conflict that was present during each trial.

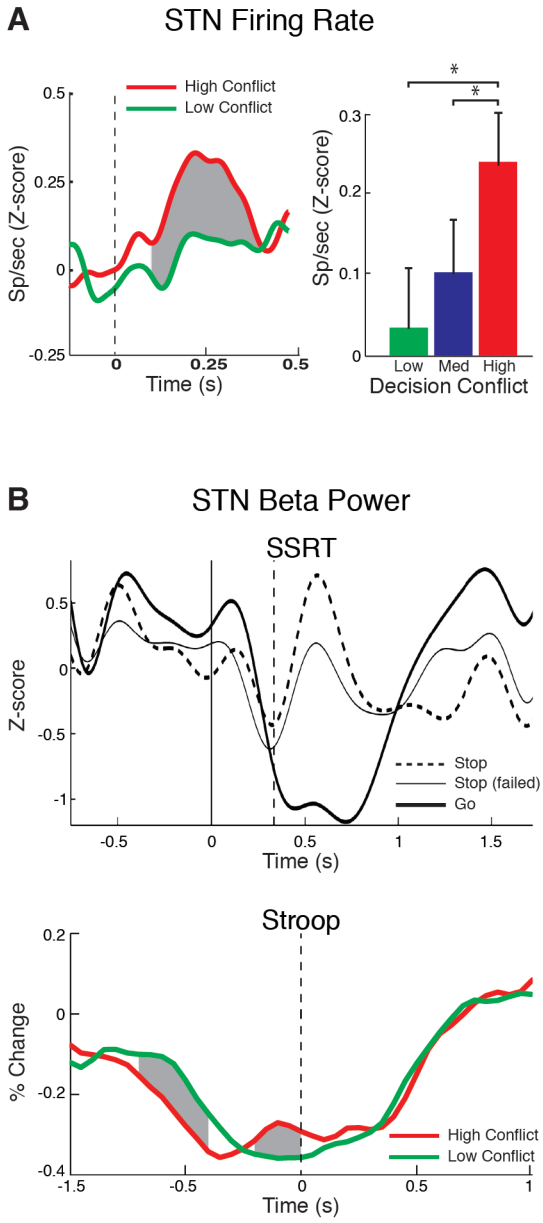


Figure 1.2. Human electrophysiology supporting a role for the STN in decision-making.
(A) STN neuronal firing. STN neurons fire more action potentials during high conflict trials in a probabilistic reward task, and degree of firing correlates with the level of conflict. Average across 27 cells recorded from 14 patients. Vertical dotted line indicates stimulus onset. Sp/sec is neuronal spikes/second. **(B)** STN LFP Beta band activity. Top: Stop Signal Reaction Time Task. STN beta power increases during abrupt stopping and is higher for correctly inhibited trials relative to failed stop trials. Average of 9 patients. Vertical solid line indicates go cue onset, and vertical dotted line indicates average stop signal onset. Bottom: Stroop Task. A similar STN LFP beta power increase occurs when subjects (n=12) perform high conflict trials (e.g. font-color-text mismatch), but not when they perform low conflict trials that do not contain incongruent information in the stimulus. Vertical dotted line indicates response. Gray shaded regions represent times where paired t tests demonstrate significant difference between congruent and incongruent traces. Note, the early differences in beta power (from -700 to -400 ms) were due to differences in the timing of when the stimulus was displayed relative to the response and do not reflect conflict related activity. Reproduced with permission from (Zaghloul et al., 2012)(A) and (Brittain et al., 2012; Ray et al., 2012)(B)).

1.4.2 Synchronized oscillations in the beta frequency (13-30 Hz) band

The single unit studies described in the preceding section, together with the similar effects of STN lesions and STN DBS on impulsivity are in line with the predictions made by the hold your horses computational model, which was largely based on the rate model of basal ganglia architecture (see section 1.1.2 "The Basal Ganglia Rate Model: The direct and indirect pathways"). As we discussed above, however (see section 1.1.4 "Beyond the rate model"), more recent studies of basal ganglia physiology have suggested that patterned changes in network-wide oscillatory basal ganglia activity might be at least as important as the raw firing rate changes of individual neurons. As much of the electrophysiological literature on the STN focuses on the potential anti-kinetic role of beta oscillations in motor circuits (see section 1.1.5 "Pathological oscillations in the basal ganglia"), several studies have sought whether beta oscillations might play a similar role in response inhibition and conflict. Kühn et al. (2004) demonstrated that during a Go-NoGo task, Go trials showed the well-established movement related decrease in beta band activity. NoGo trials likewise showed an initial decrease in beta activity; however, this decrease was quickly reversed to a relative beta increase, posited to inhibit movement. Similarly, during the stop signal reaction time task, stop trials are associated with an increase in beta band activity, with the latency of the beta increase correlating with how quickly subjects react to the stop signal (Ray et al., 2009) (Figure 1.2B). Alegre et al. (2013) have recently extended this finding to show that beta activity is higher during successful stop trials relative to failed stop trials and that elevated STN beta power is associated with elevated beta band coherence with the motor cortex. These studies all suggest that beta oscillations in the STN and connected structures may play an important role in inhibiting movement, particularly when actions should be prematurely halted.

In order to explore if beta activity is also implicated in a more cognitive task that requires responses during conflict, Brittain et al. (2012) used the Stroop task to show that high conflict trials exhibit a similar beta band pattern as observed during stopping tasks. Both low and high conflict trials demonstrated a decrease in beta band activity following the stimulus onset. During

high conflict trials, however, a small relative increase in beta power took place before the response was made (Figure 1.2B). Notably, during error trials the beta band increase still appeared, but it occurred after the response was made. In line with the race model proposed by Schmidt et al. (2013, see section 1.4.1 "Firing rate studies of single neurons" above) the authors interpreted this finding as evidence that when the beta signal arrived too late to inhibit movement, errors occurred.

So far we have considered the role of beta activity in what may be termed reactive response inhibition, in that the movement inhibition follows an imperative cue. However, beta activity may also play a role in pro-active inhibition even during tasks that do not require any movement whatsoever. This was recently investigated in a study in which PD subjects were asked to increase the count of an internally maintained number only when a trial contained a particular visual stimulus (Oswal et al., 2012). When the subjects were given an informative warning cue that let the subjects know that the upcoming trial would not contain the relevant stimulus, beta power increased and remained elevated through that trial. This result implicates STN activity in pro-active inhibition, as the increase in beta activity preceded stimuli made irrelevant by the specific warning cue. It also implicates the STN in non-motor inhibition in which an update to working memory is inhibited. Importantly, the effect was absent when the patients performed the task after their dopaminergic medication was withheld. Imaging evidence also points to the involvement of the STN in pro-active, as well as reactive, response inhibition (Ballanger et al., 2009; Castrioto et al., 2014).

1.4.3 Synchronized oscillations in the theta frequency (4-8 Hz) band

Although the firing rate and beta oscillation studies discussed thus far implicate the involvement of the STN in response inhibition, much of the literature dealing explicitly with decision-making and conflict in healthy subjects focuses on the role of oscillations of lower frequency (≤8 Hz) in the frontal cortex (Cavanagh et al., 2012). These are generally termed theta activity, although they

can extend to sub-theta (<4 Hz) frequencies as well. During an Eriksen flanker task in which subjects must indicate the direction of a middle arrow flanked by either low (<<<<<) or high conflict (>><>>) arrows, theta power increases over the frontal cortex correlate with the prolonged reaction time of high conflict trials (Cohen and Cavanagh, 2011). In a subsequent study, the same authors used the probabilistic paired choice paradigm and drift diffusion modelling to show that medial pre-frontal cortex (mPFC) theta power correlates with the amount of "evidence" needed to make a decision during conflict (Cavanagh et al., 2011). Notably, DBS of the STN completely reversed this relationship. These findings are consistent with the general predictions made by the "hold your horses" computational model in that disruption of the STN's activity leads to more errors. However, rather than assume that the STN calculates conflict based on information it receives from pre-motor cortex, the authors suggested that structures within the mPFC, such as the anterior cingulate cortex (ACC) or the pre-supplementary motor area (pre-SMA), detect conflict and subsequently communicate with the STN through theta activity to inhibit motor responses. Accordingly, they found higher low frequency (2-5 Hz) power in the STN during high conflict trials. These results indirectly suggest that theta band activity may allow the frontal cortex, which sends projections to the STN via the hyperdirect pathway (Haynes and Haber, 2013), to rapidly inhibit movement by activating the STN.

1.5 Thesis objectives

In the work described below, I attempt to explore the mechanisms by which the STN may carry out the role it is thought to play in decision-making under conflict. Towards this goal, I aim to:

1. Explore whether electrophysiological activities such as STN theta power and phase play a role in conflict processes (Chapter 3).

2. Investigate whether electrophysiological activities from the frontal cortex and the STN are coherent and if so how they relate to conflict (Chapter 4 and Chapter 6).

3. Investigate how frequency specific oscillatory activity affects the firing dynamics of the STN and how these interactions are related to conflict and action selection (Chapter 5).

4. Explore whether the STN and its coherence with cortical structures is also involved in other cognitive processes such as across trial adaptations to conflict and errors (Chapter 6).

Chapter 2. General Materials and Methods

2.1 Test Subjects and Recordings

2.1.1 Ethics

All subjects that participated in the studies presented here gave their written informed consent to take part in the experiments. None of the subjects received any form of financial compensation for their participation. All experimental protocols used were approved by the appropriate ethics committees: the National Hospital for Neurology & Neurosurgery and Institute of Neurology Joint Research Ethics Committee, London UK; the National Research Ethics Committee A, Oxford, UK; and the National Institutes of Health Institutional Review Board, Bethesda, MD, USA.

2.1.2 Post-operative recordings in the UK

All studies conducted in the UK involved post-operative recordings from the STN of PD patients who had received DBS surgery as a prelude to high frequency stimulation for the treatment of advanced PD. Techniques to target and implant electrodes in the STN have previously been described (Foltynie and Hariz, 2010). Briefly, the target for the lower-most contact was below or in the lower border of the STN and the target at the upper most contact was the upper most border of the STN. Lead location was confirmed with intraoperative stereotactic MRI at the University College London Hospital and immediate postoperative stereotactic computed tomography at the John Radcliffe Hospital and King's College Hospital. Effective stimulation was confirmed intraoperatively in patients operated at the John Radcliffe Hospital and King's College Hospital. The permanent quadripolar electrode used was model 3389 (Medtronic) featuring four platinum-iridium cylindrical surfaces with a centre-to-centre separation of 2mm. Electrode extension cables were externalized through the scalp to enable recordings before connection to a subcutaneous

DBS pacemaker, implanted in a second operation up to 7 d later. Clinical details of the patients are available in Table 3.1, Table 4.1, and Table 6.1. All recordings included in the analysis of this data were made 3-6 days after DBS electrode implantation surgery while patients were waiting to have their second surgery for DBS pacemaker implantation. These recordings were made while patients were in an ON clinical state having recently taken their Parkinsonian medication. This was done so as to reproduce "normal" basal ganglia function to the greatest extent possible. All signals were sampled at 2048 Hz, band pass filtered between 0.5 and 500 Hz, and amplified using a TMSi porti and its respective software (TMS International, Netherlands).

2.1.3 Intraoperative recordings at the NIH

All electrophysiology studies at the NIH were conducted in the intra-operative environment while patients were awake, off medication (stopped 12 h preoperatively), and engaged with either the surgical team or the experimental task. Intraoperative microelectrode recordings were used to identify the STN based on firing rate and pattern (Gross et al., 2006). Recordings were made using targeting electrodes comprised of a pair of macro and microelectrode contacts (Alpha Omega, Alpharetta, GA). Macroelectrode contacts were positioned 3 mm dorsal to the microelectrode tip. Three targeting electrodes, separately spaced 2 mm apart, were simultaneously advanced during each recording session. Raw signals were sampled at 1.5024 and 24.0345 kHz from macro and microelectrode contacts, respectively, and stored using a MicroGuide Pro data acquisition system (Alpha Omega Co., Alpharetta, GA). Mean coordinates of the microelectrode recording sites during the behavioural paradigm, referenced to the mid-commissural point, were x = 11.1 ± 0.2, y = 3.6 ± 0.5, and z = 4.9 ± 0.6 for left electrode recordings, and x = -11.8 ± 0.2, y = 3.2 ± 0.6, and z = 4.5 ± 0.6 for right electrode recordings. These coordinates demonstrated the characteristics of STN as determined by microelectrode recordings (Gross et al., 2006) and response to intra-operative stimulation and correspond to left and right STN on the Schaltenbrand-Wahren brain atlas. Recording sites during the behavioural

task were 2.6 ± 0.2 mm (56.4 ± 3.6 % of STN length) ventral to the dorsal border of the STN, as determined by the clinical microelectrode recordings. Three months after the surgery, patients returned to the clinic to repeat their experimental testing with their DBS stimulators turned on and off. No electrophysiological recordings were made in these post-operative testing sessions, and only reaction time and error rate information was stored.

2.2 Electrophysiology data analysis

2.2.1 Data pre-processing and conversion to the frequency domain

LFPs were collected with a common average reference. Prior to any analysis, all LFP signals were band bass filtered between .5 and 500 Hz, notch filtered at the mains line frequency (50 Hz for data collected in the UK and 60 Hz for data collected at the NIH), and down-sampled to 1000 Hz. All UK LFP signals were re-referenced using a bipolar montage whereby the difference was found between adjacent contacts within each STN electrode. As there were 4 vertically aligned contacts per electrode (numbered 0-3, with 0 being the ventral most contact and 3 being the dorsal most contact), this resulted in 3 bipolar contacts per STN: 0-1, 1-2, 2-3. Occasionally one contact was not active in the STN electrode, which resulted in only 1 or 2 bipolar contacts for that electrode depending on whether the inactive contact was one of the middle two (1 or 2) or outer two (0 or 3) contacts, respectively. A bipolar montage was preferred to limit the effects of volume conduction from distant sources.

All NIH data were re-referenced using a global average montage. In this procedure, the mean voltage time series averaged across the three macroelectrode LFP contacts was calculated, and the re-referenced values for each contact were found by taking the voltage difference between that contact and the global average at each time point in the recording. The global average referencing scheme was used for the following reason. Unlike the Oxford contacts that were all in a straight line in the same electrode, the NIH LFP contacts (number 0-2) were all on separate electrodes

that were parallel to each other (see Figure 5.1 for visual). Thus a bipolar configuration would result in 2 bipolar arrangements being possible for each contact (0-2, 0-1, 2-1). Given that we wanted to relate the spiking activity recording in an electrode's microelectrode contact to that electrode's macroelectrode contact, it would be unclear which of the two macroelectrode bipolars should be used for each microelectrode contact. Thus we chose to conduct the NIH analysis using a global average configuration. Repeating the analysis of the macro electrode power and phase changes using a bipolar montage, however, produced qualitatively very similar results as those presented in this thesis.

As the primary focus of this thesis was to analyse the frequency specific patterns of the STN and related cortical structures, several spectral analysis techniques were used to convert the recorded voltage traces from the time domain into the frequency domain. The first of these techniques involved the wavelet transform. The wavelet transform consists of convolving the pre-processed (see previous two paragraphs) LFP signal with complex valued Morlet wavelets. The Morlet wavelets are constructed by multiplying a sinusoid signal oscillating at the frequency band of interest with a Gaussian kernel. Convolving the LFP signal with the complex wavelet results in a time evolving instantaneous estimate of the complex valued amplitude and phase of the LFP signal at the frequency that was used to generate the wavelet. The instantaneous power can then be found by squaring the amplitude. By using several frequencies to generate a family of wavelets, the oscillatory power and phase can be found for the different frequencies of interest. For all of the analyses presented here, we used wavelets consisting of six oscillations (i.e. wave number=6). The frequency of the 47 different wavelets that were used ranged from 2 to 107 Hz, logarithmically spaced at 8 scales per octave (i.e. sampling 8 frequencies between 2 and 4 Hz, between 4 and 8 Hz, between 8 and 16 Hz, etc.). As all of the tasks that were used involved an event related organization, the continuously recorded electrophysiology signals were separated into trials, and each trial's signal was then convolved with the wavelet. Any trial that showed an

artefact in the power spectrum was discarded in an unbiased manner. This was achieved by plotting the power time series in the theta and gamma ranges and discarding any trial with an abnormally high power level. Typically the threshold for rejecting a trial was around 20 standard deviations higher than the average power observed over all trials. The time windows used for each trial varied by task (detailed in the chapters below), but typically ranged between 1-2 s before an event (such as a stimulus presentation or a response onset) occurred and 1-2 s after the event. In addition an extra 1-2 s of buffer time was used to account for any edge effects that would occur when the wavelet was convolved with zeros at the end of the time series of interest. The buffering time points were not included in the analysis.

When there was an a priori reason to focus on a specific frequency range of interest (such as from 4-8 Hz), the raw LFP trace was band-pass filtered in the frequency range on interest, and the Hilbert transform was used to generate an instantaneous estimate of the amplitude and phase of the entire signal included in the band pass filter. The benefit of using the Hilbert transform over the wavelet transform is that the loss of frequency resolution that comes with band pass filtering the signal over a wide frequency range is compensated by increased temporal resolution. Furthermore different subjects can demonstrate peak amplitudes at slightly different frequencies (Cohen, 2014), which may be diluted when the wavelet derived signals are averaged across subjects. With the Hilbert transform, the peak activity in the frequency range of interest largely determines the estimated amplitude for that signal.

2.2.2 Oscillatory power analysis

In order to generate a time-by-frequency matrix ($AveM_{time,freq}$) of the average STN activity during certain trial type (such as during high conflict trials), the following procedure was used (summarized in Figure 2.1). First (step 0), the wavelet (or Hilbert) transform was applied to all three re-referenced signals recorded from each STN. This resulted in a matrix containing the power data for each channel and each trial ($M_{channel,trial,time,freq}$). If the evoked activity was not

subtracted from the signal (see next paragraph for a description of the methods applied when the evoked activity was subtracted), the following procedure was used. First (step 1), a baseline period was chosen to which the data were normalized. Typically this corresponded to an inter-trial time period in the task in which there was very little task related activity present in the STN (such as the 1000 ms leading up to the presentation of a warning cue). Occasionally the entire trial period was used as the baseline (chapter 4, and chapter 5), although this never changed the results presented, as most of the results focused on relative differences between types of trials (such as low vs. high conflict trials). Once a suitable time period for the baseline was chosen, the average power value recorded over the entire baseline period was then found for each trial, resulting in a baseline power matrix ($B_{channel,trial,freq}$). The mean and standard deviation across trials was then found (step 2), resulting in channel-by-frequency matrices for the mean ($AveB_{channel,freq}$) and standard deviation ($STDB_{channel,freq}$) of the baseline period. The unnormalized data ($M_{channel,trial,time,freq}$ generated in step 0) was then normalized by subtracting the mean baseline matrix ($AveB_{channel,freq}$) and dividing by either the mean baseline matrix ($AveB_{channel,freq}$) for a measure of the percentage change from baseline or the standard deviation baseline matrix ($StdB_{channel,freq}$) for a z-scored measurement of the changes from baseline (step 3). This resulted in a matrix of normalized power data ($NormM_{channel,trial,time,freq}$). The resulting matrix was then averaged across trials (step 4) and across channels (step 5) to generate a time-by-frequency matrix for each STN ($NormM_{time,freq}$). Finally, the STN matrices were averaged across STNs (step 6) to generate the average normalized matrix ($AveM_{time,freq}$). Steps 4-6 were done separately for low and high conflict trials (or any other subgroups of trials being compared i.e. correct high conflict trials vs. incorrect high conflict trials). In order to test for differences between low and high conflict trials, the difference between the low conflict matrix ($AveM_{time,freq}$_low) and the high conflict matrix ($AveM_{time,freq}$_high) was found resulting in difference matrix ($D_{time,freq}$, step 7). This matrix was then compared to 1000 matrices of permuted data, each of which was generated by shuffling the low and high conflict trials generated in step 3 before averaging across trials in

step 4. By comparing the real difference matrix to the 1000 permuted matrices, a p-value could be assigned to each time-frequency point in the difference matrix. Exceedence mass testing (described below in section 2.3.2 ”Permutations and exceedance mass testing”) was then used to correct for multiple comparisons (step 8). When we analysed data using the Hilbert transform (see preceding paragraph), the same procedure was used, with the only exception being that the frequency dimension of all of the matrices described in this paragraph was only one unit long.

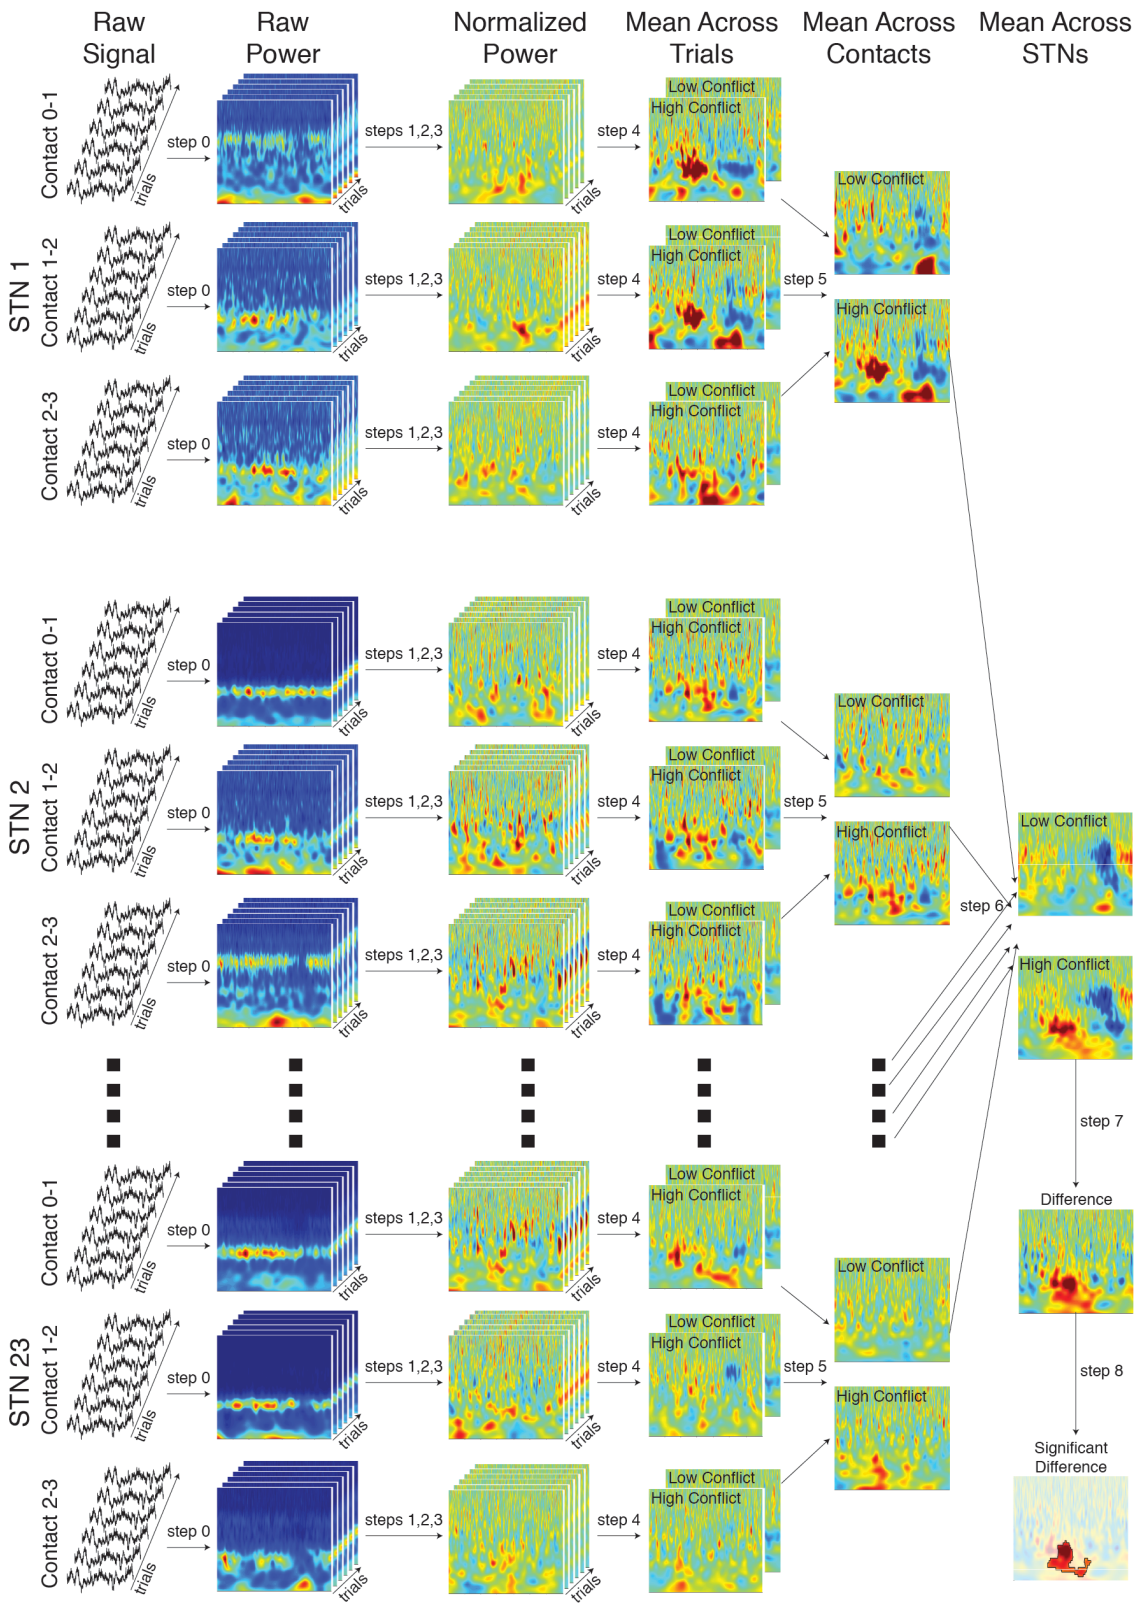


Figure 2.1. Electrophysiology data processing for power analysis.
See text for description of each step. Schematic shown is for a hypothetical dataset with 23 STNs.

Occasionally (chapter 4, and chapter 5), we were interested in whether power differences between low and high conflict trials reflected induced or evoked activity. Induced activity involves changes in power that arise from an increase in amplitude at the frequency of interest. Evoked activity involves changes that arise from an increase in inter-trial phase consistency (described below). In order to isolate the power differences arising from induced and evoked activity, the following procedure was applied to the data resulting from step 0 of the preceding paragraph. First, the average evoked power was found for each channel by averaging that channel's LFP time series data across trials. The wavelet transform was then applied to the resulting matrix, affording a matrix of evoked power ($E_{channel,time,freq}$) from each STN. This matrix was then subtracted from the power data matrix generated in step 0 of the preceding paragraph ($M_{channel,trial,time,freq}$). Averaging the resulting matrix across trials led to a matrix of induced power data ($I_{channel.time,freq}$). The baseline for each frequency ($B_{channel,freq}$) was found by averaging the matrix across the baseline time points. This baseline matrix was then subtracted from each time-frequency point of $I_{channel.time,freq}$. Steps 5-8 of the preceding paragraph were then applied. The resulting matrix was averaged across channels to generate a normalized matrix for each STN (step 5). The resulting matrix was then averaged across STNs (step 6) and the difference between the low and high conflict trial matrices was then found (step 7) and compared to 1000 surrogate difference matrices (step 8) found by permuting the low and high conflict trials prior to the calculation of the evoked power.

Comparing the induced and evoked activity high – low conflict differences to the induced only high – low conflict differences revealed very similar differences suggesting that the power difference matrices are dominated by induced activity and thus the subtraction of the evoked activity was for the most part unnecessary. Nevertheless, as our first study focused on differences in power and ITPC, we wanted to ensure that any differences in one metric were not secondary to differences in the other metric. As this was found not to be the case, only the first two studies

presented here were analysed by subtracting the evoked activity. The remaining studies were analysed using only the steps described in the first paragraph of this section as outlined in Figure 2.1.

2.2.3 Oscillatory Inter-trial phase consistency (ITPC)

Figure 2.2A shows the instantaneous phase of several trials aligned to the presentation of a stimulus. The onset of the stimulus induces a phase reset of the on-going oscillation, resulting in a similar phase being recorded at the same time point of each trial. In contrast, the traces shown in Figure 2.2C do not have phase consistency across trials. Inter-trial phase consistency (ITPC, also known as "inter-trial phase clustering" (Cohen and Gulbinaite, 2014)) is a way of quantifying how similar the phases are across trials at a particular time point in an experiment. Applying the wavelet (or Hilbert) transform to LFP voltage traces results in a matrix of instantaneous data (step 0 in the preceding section 2.2.2). The phases at a given time point can then be projected on the complex plane for all trials, and the consistency of those phases can be found by averaging across trials and taking the absolute value:

$$\text{ITPC(t)} = \left| \frac{1}{n} \times \sum_{t=1}^{n} e^{i\theta_t} \right|$$

Equation 2.1. Inter trial phase consistency (ITPC).

where n is the total number of trials, and θ_t is the phase angle at time t. Using this formulation, an ITPC(t) value of 0 would mean there is a uniform distribution of phase across trials at time t, and a value of 1 would mean that the phase at time t is identical for each trial. Figure 2.2B shows the ITPC values over time for the data shown in Figure 2.2A.

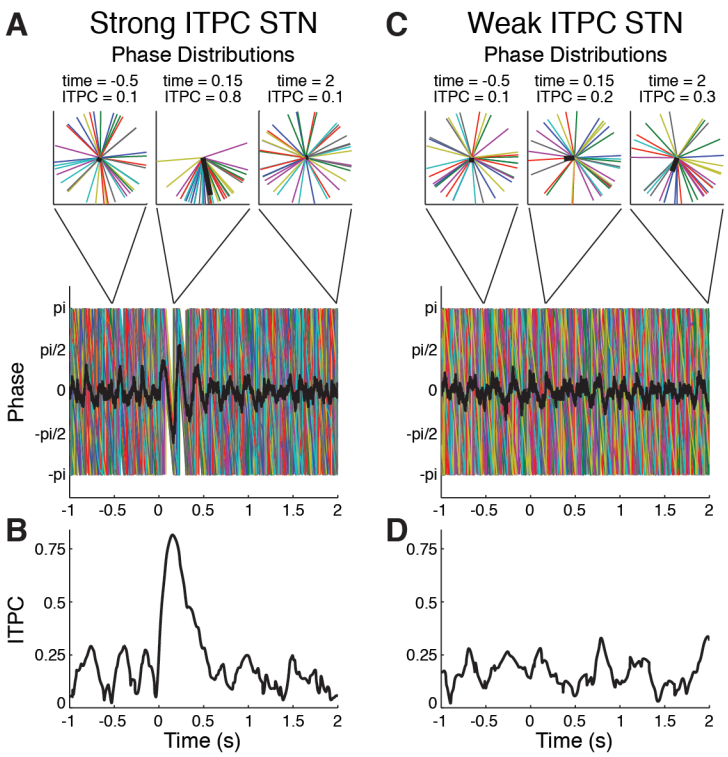


Figure 2.2. Inter-trial phase consistency (ITPC).
(A) Bottom plot shows the instantaneous theta phase (4-8 Hz band pass filtered and extracted with the Hilbert transform) for all of the low conflict trials in one bipolar contact in one STN. Solid black line shows the average instantaneous phase at each time point. Time=0 corresponds to the presentation of a stimulus that resets the phase of the theta oscillations in this recording. Top plots show the phase distribution at three time points throughout the recording (t=-0.5, t=0.15, and t=2). The average vector of each distribution of phase vectors is indicated by the thick black line in each plot. The length of this vector is what was used to quantify the inter-trial phase consistency (ITPC). The distributions at t=-0.5 and t=2 are very close to uniform and thus have an average vector length that is very close to 0 (ITPC=0.1). The distribution at t=0.15 is less uniform and thus has an average vector length that is closer to 1 (ITPC=0.8). **(B)** The ITPC for each time point in the recording shown in A. **(C-D)** Same as (A-B) but for a recording that did not show a clear increases in ITPC when the stimulus was presented.

ITPC values can be heavily biased by the number of trials included in the analysis (as an extreme example, the ITPC of 1 trial would be 1). Therefore, all comparisons between low and high conflict trials (or between any other subsets of trials), involved 1. down-sampling the largest trial group to match the number of trials in the smaller trial group, and 2. calculating the ITPC of the down-sampled data. This was repeated 1000 times and the resulting 1000 down-sampled estimates of ITPC were averaged to get the best estimate of the down-sampled ITPC. As an example consider a data set in which there were 80 high conflict trials and 40 low conflict trials (due to a 2:1 ratio of high to low conflict trials, see Chapter 3 and Chapter 6). The ITPC for the 40 low conflict trials was found using the procedure described in the preceding paragraph. In

order to estimate the ITPC for the high conflict trials, however, 40 of the 80 high conflict trials would be randomly chosen for the first estimate of the high conflict trial phase consistency ($ITPC_{high1}$). This procedure would be repeated 1000 trials resulting 1000 estimates of the high conflict trial ITPC ($ITPC_{high1:high1000}$). The resulting $ITPC_{high}$ values would then be average to generate $ITPC_{high_downsampled}$ and this value would be compared to $ITPC_{low}$.

In most cases, the raw ITPC values were compared between conditions as the down-sampling of trials accounted for any biases in trial counts that might affect the low vs. high conflict trial comparison, and the ITPC value range was restricted from 0 to 1 by definition. For Chapter 6, however, we normalized the ITPC values, so all of the metrics presented in Figure 6.2 and Figure 6.3 would be on the same percentage change scale. This process involved finding each channel's average ITPC value at baseline for a down-sampled number of trials (chosen to match the smallest number of trials used in the comparison, 40 in the exampled used in the preceding paragraph), and comparing each time point's ITPC value to that baseline.

2.2.4 Inter-site phase coherence

Similar to ITPC, inter-site phase coherence involves finding the consistency of phase values across time. In the case of phase coherence, however, the phase values used are those found by subtracting the instantaneous phase at one recording from the instantaneous phase at the same time, but in a different recording. Thus phase coherence measures how consistent the phase difference is between two locations across trials:

$$\text{phase coherence(t)} = \left| \frac{1}{n} \times \sum_{t=1}^{n} e^{i\theta_t} \right|$$

Equation 2.2. Phase coherence

where n is the total number of trials, and θ_t is the phase difference ($\theta_{LFP} - \theta_{EEG}$) angle at time t. Using this formulation, a phase_coherence(t) value of 0 would mean that the two signals have

perfectly independent phases at time t across trials, and a value of 1 would mean that the phases have an identical phase difference at time t for each trial. Though the mathematical formula used to calculate phase coherence is similar to that used to calculate ITPC, it is important to emphasize that ITPC measures the consistency of the *absolute* phase in *one* signal and phase coherence measures the consistency of the phase *difference* between *two* signals. High phase coherence typically indicates a high level of synchrony between the two recordings as their phases reflect a consistent time lag between the two sites across trials. Similar to ITPC, phase coherence can be affected by the number of trials used to calculate the measurement, so all comparisons between low and high conflict trials involved a down-sampled data set of the larger trial group that was chosen to match the number of trials in the smaller trial group (see second paragraph of previous section 2.2.3 "Oscillatory Inter-trial phase consistency (ITPC)" above). Like our ITPC calculations, down-sampling was done 1000 times.

Unlike the raw phase values that are used to calculate ITPC, the phase difference value used to calculate inter-site phase coherence at time t can and often will be very similar at time t-2 ms, t-1 ms, t, t+1 ms, t+2 ms, t+3 ms, etc. This is because the intersite phase coherence measures the consistency of the phase difference between two signals, which will not vary substantially across short periods of time due to the oscillatory nature of both signals (ITPC, on the other hand, measures the consistency of the raw phase at one site, which will vary across time by the very definition of being an oscillation). If the two signals are independently oscillating (i.e. they have a low level of synchrony), the phase difference at time t will be similar at time t+1 but will typically be different at time t+100 due to random fluctuations in either of the two signals. When there is high level of synchrony between two signals, however, there is an increased likelihood that the phase difference at time t+100 will be similar to the phase difference at time t because the two oscillations are fluctuating together (Lachaux et al., 2002). For this reason, whenever we analysed phase coherence between the STN and scalp electrodes, we looked at how consistent the

phase difference was across trials (just like ITPC) and across time (unlike ITPC). This procedure is outlined in detail in the following section 2.2.5 “Inter-site Lachaux coherence”.

All phase coherence values presented here were normalized to the phase coherence values recorded during a “baseline” period. This process involved finding each pair of channel’s average phase coherence value at baseline (using the same number of trials as the smallest trial group in the comparison being made), and comparing each time point’s phase coherence value to that baseline.

2.2.5 Inter-site Lachaux coherence

Another metric for measuring inter-site connectivity that is highly related to phase coherence is the frequency domain cross-correlation, most commonly referred to simply as coherence (Gardner, 1992; Pesaran, 2008). It is similar to phase coherence in that it measures the consistency of the phase difference between two signals at a given frequency, but it also takes into account the correlation of the two signal’s power values at that frequency. Coherence is estimated by calculating the average cross spectrum of the two signals and dividing by the product of each signal’s auto-spectrum:

$$\varrho\,(f) = \frac{|S_{xy}(f)|}{[S_{xx}(f) \,.S_{yy}(f)]^{1/2}}$$

where

$$S_{xy}(f) = \sum_{i=1}^{N} x_i(f).y_i^{*}(f)$$

Equation 2.3. Traditional coherence.

$x_i(f)$ and $y_i(f)$ are the discrete Fourier coefficients at frequency f of the finite time series x(t) and y(t), and N are different time windows from which x_i and y_i are estimated. The N time windows can be calculated by dividing a continuous time segment into N windows or by calculating $x_i(f)$ across N trials for the same time window in event locked data. The Granger causality and the phase coherence results presented in Chapter 4 section 4.3.3 were confirmed using this traditional coherence metric (data not shown).

More recently, a modification of the coherence metric has been proposed to analyse wavelet (and Hilbert) derived instantaneous measurements of inter-site connectivity (Lachaux et al., 2002). In contrast to the traditional Fourier based coherence shown in the equation above, the averaging over the window N is replaced by an integration across time over a window that is proportional to the frequency being analysed:

$$WCo(t,f) = \frac{|SW_{xy}(t,f)|}{[SW_{xx}(t,f) \cdot SW_{yy}(t,f)]^{1/2}}$$

where

$$SW_{xy}(t,f) = \int_{t-\delta/2}^{t+\delta/2} W_x(\tau,f) \cdot W_y^*(\tau,f)\, d\tau$$

Equation 2.4. Lachaux coherence.

W_x and W_y are the wavelet (or Hilbert) transform of the signals x(t) and y(t), and $\boldsymbol{\delta}$ is a scalar that depends on frequency. By integrating across time and across trials for our data, we were able to generate a time evolving estimate of the coherence between the STN and the frontal cortex.

Low and high conflict trials were analysed separately. Prior to any analysis, the raw data were band passed filtered in the theta band and the Hilbert transform was used to extract the power and

phase of each LFP bipolar channel and the Fz-Cz bipolar channel. The cross spectrum at each time point was then found by multiplying the LFP's complex Hilbert value at each time point in each trial by the conjugate of the EEG's complex Hilbert value for the corresponding time point. The complex cross spectrum values were then averaged across trials and a sliding window was used to integrate across time. The width of the window ($\boldsymbol{\delta}$) was chosen to be 1.6 s for the dots task (Chapter 4) and .333 s for the flanker task (Chapter 6). Different windows were chosen based on the experimental design of the task (the reaction times for the dots and flanker tasks were around 3.4 s and 0.5 s, respectively). The absolute value of the resulting average cross spectrum was then divided by the product of the two signals' autospectum to produce the Hilbert coherence. The Hilbert coherence between the LFP bipolar channels and the Fz-Cz bipolar channel was found for each LFP channel separately and averaged across channels to generate one time-by-frequency Lachaux coherence matrix for each STN for each condition being compared (for example high and low conflict trials). The same across-time integration techniques were applied to the inter-site phase coherence metric (described in the preceding section) using only the phase values. This was achieved by replacing all amplitude values in the Lachaux coherence formula above with the number 1. Whereas the simple phase coherence at time t described in the preceding section (section 2.2.4 "Inter-site phase coherence") involves looking at the average of the phase difference at time t across all trials, the Lachaux phase coherence at time t involves looking at the average of the phase difference at time t-window_length/2 to t+window_length/2 across all trials.

2.2.6 Granger Causality analysis

One of the primary interests of this thesis was to explore the directionality of cortico-STN electrical coupling during conflict (chapter 4). In order to do so, we utilized multivariate autoregressive (MVAR) modelling to estimate the coherence (Gardner, 1992; Pesaran, 2008) and Granger causality (Granger, 1969; Kamiński and Blinowska, 1991; Ding et al., 2006; Seth, 2010)

between LFP signals recorded from the STN and EEG signals recorded over the frontal and posterior cortices. When I say that one signal is "Granger-causal" to a second signal, I mean that the inclusion of past observations of the first signal reduces the prediction error of the second signal in a linear regression model (relative to a model that only includes past observations of the second signal; (Granger, 1969; Seth, 2010)). In order to pre-process the data and calculate the MVAR estimated coherence and Granger causality, the Granger causality connectivity analysis (GCCA) toolbox (Seth, 2010) was used to analyse the raw, bipolar EEG and LFP signals.

Prior to any analysis, 50 Hz mains artefact was removed by notch filtering, and the data were down-sampled to 250 Hz. This was done so as to improve frequency resolution while maintaining the time scale of the interactions (Schlögl and Supp, 2006; Ruiz et al., 2013). Because our recordings took place over multiple trial iterations (as opposed to one long recording) each trial was treated as an independent realization of a statistically stationary process. In order to minimize non-stationarities across trials, each bipolar channel was normalized horizontally and vertically as follows (Ding et al., 2000; Seth, 2010). First, each trial was zero-meaned across time (horizontal normalization) by subtracting the mean voltage amplitude of that trial from each time point in the trial and dividing the resulting values by the standard deviation of the voltage amplitude of that trial (Ding et al., 2000). Next, the mean and standard deviation of the evoked response (averaged across all trials) was calculated and each corresponding time point in each trial was z-scored (vertical normalization) by the appropriate values corresponding to that time point in the evoked response (Ding et al., 2000; Seth, 2010). The vertical normalization step was done separately for the low and high conflict trials. Though vertical normalization is consistently used in the literature (Brovelli et al., 2004; Ruiz et al., 2013) in order to "(1) remove the first-order nonstationarity from the data and (2) make the ensemble mean equal to zero"(Ding et al., 2000), we found qualitatively very similar results when the vertical normalization step was omitted (data not shown). This is most likely due to our use of gradual changes in the dot coherence that never

resulted in a stimulus-onset triggered evoked response. After the data were pre-processed, each 6 s trial was windowed into thirteen 1000 ms blocks with a 500 ms overlap between blocks.

The appropriate model order (i.e. the number of time lags to include when generating the multivariate autoregressive (MVAR) model matrix) was then found for each STN:EEG electrode pair as follows. First, the best model order of each bipolar LFP:EEG pair was found for each of the conflict trial time windows using the Bayesian information criterion (BIC, (Schwarz, 1978)). This resulted in 39 model orders for each STN:EEG pair (13 time windows by 3 bipolar LFP channels); these 39 model orders were then condensed into one number by averaging across the 13 time windows and the three LFP contacts in that STN. Therefore, each STN finished with two model orders: one for the STN:Fz-Cz interaction and one for the STN:Pz-Cz interaction. There were no significant differences between the STN:Fz-Cz and STN:Pz-Cz model orders across subjects (8.74 vs. 8.79; $p>0.05$, paired t-test). The average difference in model order between the STN:Fz-Cz analysis and the STN:Pz-Cz analysis was 0.05 orders, which is approximately equal to 0.2 ms for a 250 Hz sampling rate. Given our 250 Hz sampling rate, our model order of ≈ 9 steps corresponds to "looking into the past" for approximately 36 ms. This model order is consistent with the model orders others have used to resolve frequencies in the delta and beta bands (Brovelli et al., 2004; Schlögl and Supp, 2006; Ruiz et al., 2013). Ruiz et al., for example use a model order of 7 with a sample rate of 250 Hz to analyse Granger connectivity in frequencies as low as 4 Hz (2013). Nevertheless, we also repeated our Granger connectivity analysis using a higher model order (20) and a lower sampling rate (100 Hz), which resulted in a much finer frequency resolution as each MVAR matrix contained 200 ms of data when these parameters were used. This produced qualitatively very similar results as those obtained using the BIC model order.

Once the appropriate model order had been found for each STN:EEG pair, the coherence (Gardner, 1992) between each LFP:EEG bipolar pair was found by calculating the MVAR model matrix for each trial, averaging the MVAR matrices across trials and converting the matrix into the spectral domain using the 'cca_pwcausal' function of the GCCA toolbox (Geweke, 1982; Kamiński and Blinowska, 1991; Ding et al., 2000; Seth, 2010). The formulas used by the GCCA toolbox to convert from the MVAR matrix to the spectral coherence are outlined by Ding et al. (2006). Low and high conflict trials were analysed separately for each electrode pair using the same model order determined for that pair. In order to determine the directionality of each STN-EEG interaction, the directed transfer function (DTF) provided in the Neurophysiological Biomarker Toolbox (NBT) (http://www.nbtwiki.net/) was used to find the directed coherence for each LFP: EEG pair (Kamiński and Blinowska, 1991). The frequency range used for the coherence and DTF analysis was the same as the logarithmic range used for the power analysis described above (2 to 107Hz, 8 scales per octave).

In order to verify that the model order we used was indeed capable of analysing the frequencies we analysed (2 to 107 Hz), we also used the GCCA toolbox to analyse two pairs of simulated signals with known connectivity (Figure 2.3). The sampling rate was the same sampling rate we used in the analysis of our real data (250 Hz), and the model order was the same as the average model order used in our analysis (9 points). The left column of Figure 2.3 shows the results for two signals in which signal A Granger caused 30 and 70 Hz oscillations in signal B. Signal B did not cause any oscillations in signal A. Accordingly, there is high coherence at 30 and 70 Hz, and this coherence is due to signal A Granger causing 30 and 70 Hz oscillations in signal B. The right column shows the results for the same signals modified in the following ways: 1. SignalB_filt was produced by band-pass filtering signal B between 1 and 5 Hz. 2. SignalB_filt was then shifted by 40 ms. 3. The time-shifted, filtered signal was then added to signal A. 4. Signal B was left unchanged. Due to the time-shifted filtered signal that was added to signal A, one would

expect to see an increase in the Granger causality of signal B to signal A, but not an increase in the Granger causality of signal A to signal B. Accordingly, the analysis on the right column shows an increase in delta band coherence and Granger causality from signal B to signal A. The Granger causality from signal A to signal B remained unchanged, as did the coherence in the higher frequencies. Further validation of our methods stems from that fact that very similar results (higher 2 to 8 Hz STN:Fz-Cz coherence during conflict) were obtained when we used the Fast Fourier transform (see (Pesaran, 2008)) or the wavelet transform (see previous section above (Lachaux et al., 2002)) to calculate the coherence between the cortex and the STN.

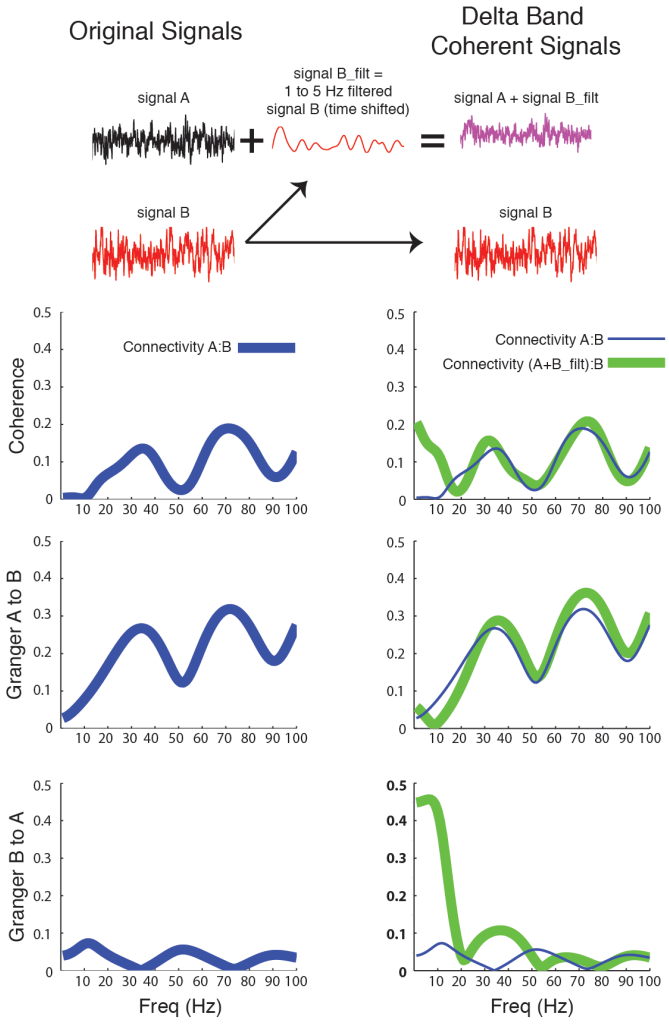


Figure 2.3. Granger connectivity analysis.
The GCCA toolbox (Seth, 2010) was used to analyze the Granger causality between LFP and EEG recordings made while patients performed a decision-making task. Along with the analysis of the real data, the toolbox was also used to

analyze two simulated signals of known connectivity where one signal (signal A) Granger caused 30 and 70 Hz oscillations in signal B. The left column, which shows the coherence and Granger causality analysis of signals A and B, shows high coherence and A to B Granger causality for these two signals in the 30 and 70 Hz bands. The right column shows the results for the same analysis when signal A was manipulated in a way that resulted in signal B Granger causing 1 to 5 Hz oscillations in signal A (to this end signal B was filtered, time shifted by 40 ms and then added to signal A, while signal B was left unchanged). In line with these manipulations to signal A, the right column shows elevated coherence in the lower frequencies (<10 Hz) and a corresponding elevation in Granger causality from signal B to A. The Granger causality from signal A to B remained unchanged. The results from the left column are re-plotted in the right column for ease of comparison. Reproduced with permission from (Zavala et al., 2014)

2.3 Statistical analysis

2.3.1 Parametric statistics

All statistics were conducted in Matlab. When comparing averages of reaction time, error rates, correlation coefficients (derived using Spearman's correlation), a simple t-test was used to compare low and high conflict trials (or any other sub-combination of trials). The only exception to this was when we compared the response times for the congruent dot motion experiment (Chapter 4), in which the reaction times were highly variable both within and across subjects. If more than two conditions were being tested, analysis of variance (ANOVA) techniques were used. For all comparisons, a significance threshold of $p<0.05$ was used. When several hypotheses were tested using the same comparison, multiple comparisons were corrected for using the family-wise error rate (Bonferroni) correction. This involved dividing the significance threshold by the number of comparisons (i.e. $p<0.0125$ would be the new threshold if four hypotheses were tested as in Figure 5.8).

2.3.2 Permutations and exceedance mass testing

Whenever we tested for differences along a continuum of time, frequency, or both, non-parametric permutation testing was used. These techniques were chosen because they allowed us to test several time-frequency points at once while allowing us to correct for multiple comparisons using exceedence mass testing (Maris and Oostenveld, 2007). This permutation-based procedure identifies contiguous windows where the metric of interest differs between trial types while correcting for multiple comparisons and avoiding a priori assumptions about

particular temporal regions and frequencies. It is a less rigid form of multiple comparisons correction than the family-wise error rate correction because it takes into account when two significant points are adjacent to one another in time, frequency, or both.

When testing for differences between low and high conflict trials in a particular metric (i.e. normalized power, ITPC, phase coherence, etc.), the average difference was found across all STNs. This procedure generated a time-by-frequency matrix of difference values (for example see Figure 2.1). The high and low conflict trial labels were then mixed and re-divided to generate a permuted low conflict trial group and a permuted high conflict trial group (where both trial groups contain low and high conflict trials randomly selected from the original groups). An average difference between permuted low and permuted high conflict trials was then found for the permuted data set, which generated a new time-by-frequency matrix of difference values. This procedure was repeated 1000 times to generate 1000 time-by-frequency matrices of permuted differences. The true difference at each point in the original matrix was then compared to the distribution of the 1000 matrices at the corresponding time-frequency point to generate a p-value for each point. The same was done for each of the permuted matrices. A threshold value was then chosen ($p<0.05$ in all of the analysis presented here). Then for the real difference matrix and for each of the 1000 surrogate matrices, a cluster statistic was calculated for each set of supra-threshold contiguous points by taking the sum of the inverse of the p-values that were in that cluster. In this manner, large cluster statistics could arise either from strongly significant differences in power between trial types or from weaker differences that persist over larger areas of time and frequency. For each of the surrogate matrices, the largest observed cluster statistic was stored thus generating an empiric distribution of maximum cluster statistics that would arise by chance. To generate a p-value for each cluster observed in the true dataset, the position of the true cluster statistic was compared to the distribution of maximum cluster statistics from the

permuted cases. Clusters were determined to be significant if their p-value calculated in this manner was less than 0.05.

In many cases in which this technique was used, clusters were found by looking at both time and frequency. In cases where we had an a-priori reason to look within a particular frequency band, clusters were only computed across time. In one instance (Figure 5.8), we averaged across time and applied the cluster correction across frequencies.

Chapter 3. Subthalamic nucleus local field potential activity during the Eriksen flanker task reveals a novel role for theta phase during conflict monitoring

3.1 Introduction

Much of the literature dealing with the involvement of cortical structures in conflict monitoring and response inhibition focuses on the role of theta band (4-8 Hz) oscillations (Cavanagh et al., 2012). mPFC theta oscillatory power is higher during high conflict trials of the Stroop and Simon tasks (Cavanagh et al., 2012; Oehrn et al., 2014) and has been shown to correlate with reaction time during the Eriksen flanker task (Cohen and Cavanagh, 2011). Prior to the beginning of our experiments, a recent paper (Cavanagh et al., 2011) had contextualized the relationships between mPFC theta, conflict, and reaction times in the framework of the drift diffusion model of decision-making (Ratcliff and McKoon, 2008; Ratcliff and Frank, 2012). In this model, a noisy neuronal process accumulates information supporting one of two alternatives until there is enough evidence to cross a "decision threshold" and execute the winning response. Cavanagh et al found that during the high conflict trials, the level of theta power recorded over the frontal cortex correlated with the "decision threshold" that was found when the reaction times were fit to the drift diffusion model. Crucially, STN DBS was found to disrupt the relationship between mPFC theta power and decision threshold. The authors interpreted their results as suggesting that the mPFC uses theta band oscillations to activate the STN and temporarily inhibit all responses during conflict. During DBS, the oscillatory patterns of the STN are disrupted resulting in impaired response inhibition and impulsivity during conflict (Frank et al., 2007; Coulthard et al., 2012; Green et al., 2013). The authors bolstered their claim of theta band mediated communication between the mPFC and the STN, by recording from the STN intraoperatively while patients performed the same task. During conflict, the STN showed higher power in a low frequency band that was close to the theta band (2.5-5 Hz).

In order to test whether the STN showed higher theta power during high conflict trials of the much simpler Eriksen flanker task, we record LFP activity in the STN in PD patients implanted with DBS electrodes while the patients performed this task.

3.2 Methods

3.2.1 Subjects

All subjects gave their written informed consent to take part in the study, which was approved by the appropriate local ethical committees. Thirteen subjects (ten males, mean disease duration 10 years, mean age 60 years, range 48–73 years) underwent bilateral implantation of DBS electrodes into the STN, as a prelude to high frequency stimulation for the treatment of advanced PD. Only one patient had been diagnosed as having an impulse control disorder. Clinical details of the patients are available in Table 3.1. The mean percentage improvement in the motor section of the Unified Parkinson's Disease Rating Scale (UPDRS) on treatment with levodopa was 59 ± 5.2% ($p < 0.001$, Wilcoxon signed-rank test between ON and OFF levodopa scores) across subjects, indicating good responsiveness to levodopa in the participants used in this study. The task was performed while patients were on their regular Parkinsonian medication.

Table 3.1. Clinical details from first Flanker experiment

Case	Age	Disease Duration	UPDRS Off (III)	UPDRS On (III)	First Symptom	Reasons for Surgery	Daily Medication (mg/d)
1	48	12	56	33	Leg stiffness	Off state immobility	Levodopa 400; Ropinirole 3
2	66	3	27	22	Tremor	Tremor	
3	62	9	47	17	Rigidity	Immobility, stiffness, tremor	Levodopa 150; Pramipexole 2.1; Pramipexole 1.1; Levodopa 750
4	52	9	29	11	Weakness in left wrist	Gait difficulties	Entacapone 600; Levodopa 600; Pramipexole 1.05
5	54	5	28	6	Tremor & leg stiffness	Gait difficulties	Levodopa 600; Pramipexole 0.35
6	64	15	31	8	Slowness & tremor on right hand	Dyskinesias, unpredictable on/off fluctuations	Apomorphine 12; Levodopa 450; Amantadine 200; Rotigotine 16; Selegiline 10

7	63	8	43	10	Slowing swallowing, stiffness, pain	Off-freezing	Rasagiline 1; Levodopa 300; Apomorphine 100
8	63	14	33	6	Frozen shoulder & arm flexibility	Off state gait difficulties	Rasagiline 1; Pergolide 3; Entacapone 400; Levodopa 500
9	56	6	52	19	Left side tremor	Tremor	Entacapone 800; Levodopa 400; Rotigotine 8
10	73	14	35	15	Right side tremor	Motor fluctuations	Levodopa 700; Rotigotine 16; Selegeline 10
11	63	14	35	24	Right side tremor	Tremor	Levodopa 150; Ropinirole 2; Ropinirole 21,
12	58	10	42	20	Leg cramp	Tremor	Trihexyphenidyl 3; Levodopa 600; Rasagiline 1; Amantadine 100
13	62	10	20	8	Left side tremor & bradykinesia	On/off fluctuations, tremor & impulse control disorder	Levodopa 1000; Trihexyphenidyl 6

3.2.2 Task

Patients performed an arrow version of the flanker task (Eriksen and Eriksen, 1974) three to six days after electrode implantation. A schematic of the task is shown in Figure 3.1. Each trial began with a black screen containing a white fixation dot in the middle of the screen, which subtended a visual angle of ≈1 degree. 500 ms before the arrows were shown, the dot changed from white to grey to prepare the test subject for movement. Either congruent (>>>>>, henceforth referred to as low conflict) or incongruent (<<><<, henceforth referred to as high conflict) arrows (visual angle ≈3 degrees per arrow) were then briefly shown and replaced with the white fixation dot after 200 ms. The subjects had two more seconds in which to respond (2.2 s total possible window for a response) before the fixation dot changed from white to grey again to signal the next trial. Correct responses were indicated by a button press in the hand corresponding to the direction of the middle arrow. The ratio of high conflict trials to low conflict trials was 2:1. Subjects underwent two 60 trial blocks. The second block for one subject was discarded due to a sharp drop off in performance after the first block.

3.2.3 LFP data recording and analysis of power

All signals were sampled at 2048 Hz, band-pass filtered between 0.5 and 500 Hz, and amplified using a TMSi porti and its respective software (TMS international, Netherlands). Monopolar LFP recordings were subsequently converted off-line to a bipolar montage between adjacent contacts (three bipolar channels per side) to limit the effects of volume conduction from distant sources.

For the comparison of correct low and high conflict trials, all incorrect responses (12±2% of all trials) as well as any trials with reaction time greater than 2 s (including trials with no response) or shorter than 150 ms were discarded (<1%). From the remaining trials, any trial with a response that was greater than three standard deviations (SDs) above the mean latency of the other trials was discarded (mean final trial count = 109 ± 6.7 trials per subject). Prior to further analysis, LFP data were filtered between 1 and 500 Hz and down-sampled to 1000 Hz. The power and phase of the bipolar LFPs in each trial were calculated using the Morlet wavelet at 8 scales/octave from 2 to 107 Hz. Each trial was analysed from 0.5 s before to 2 s after flanker arrow onset for the cue-aligned analysis and from 1.5s before to 1.5 s after the response for the response-aligned analysis. A 1 s buffer on either side was used when calculating the wavelet power to eliminate any edge effects. Any trial with a clear artefact was discarded.

To assess the differences in induced power between high and low conflict trials, we used the "induced power differences" approach outlined in section 2.2.2 "Oscillatory power analysis". Rather than using a specific baseline period, each frequency was normalized to the overall mean power of that frequency across all time points in all trials. Qualitatively very similar results were obtained when the data were normalized to the mean power recorded during just the pre-warning cue period (data not shown). Finally, all three bipolar contacts were averaged together and the difference between the low and high conflict trials was calculated. Averaging across all the contact pairs in a given electrode was performed so as to avoid selection bias, although this procedure might serve to underestimate spectral changes in the STN as not all contact pairs were

necessarily within or bridging this nucleus. The above procedure was repeated for each STN's electrode to find the mean difference across subjects. To assess the statistical significance of this difference, the same procedure was repeated 1000 times with the low and high conflict labels randomly assigned during each permutation, and exceedence mass testing (Maris and Oostenveld, 2007) was used to correct for multiple comparisons (see section 2.3.2 "Permutations and exceedance mass testing" above). When performing other comparisons (i.e. low conflict vs. fast-high conflict, low conflict vs. slow-high conflict, fast-high conflict vs. slow-high conflict, or high conflict error trials vs. correct high conflict trials), the same procedure was repeated using the relevant trial groups. To generate the theta band time-frequency time series the same procedure was used, but the values across 4-8 Hz were first averaged prior to any other calculation.

3.2.4 Inter-trial phase consistency (ITPC)

To analyse ITPC, the LFP signal was band pass filtered from 4 to 8 Hz and the Hilbert transform was used to calculate the instantaneous phase of each trial. The ITPC at each time point was then found by using the procedure outlined in section 2.2.3 "Oscillatory Inter-trial phase consistency (ITPC)" above.

To assess the statistical differences between conditions, the difference between the two ITPCs was first calculated for each bipolar signal and averaged across all three bipolar contact pairs of the electrode before being averaged across subjects. The same procedure was then repeated 1000 times with the labels of the two trial groups being randomly assigned at the beginning of each permutation. The p value at each point was calculated using the distribution of the 1000 permuted values and corrected for multiple comparisons at a significance level of 0.05 using exceedence mass testing (Maris and Oostenveld, 2007), see methods section 2.3.2 "Permutations and exceedance mass testing" above. To ensure that the results we reported were specific to the 4-8 Hz band, the same procedure was repeated using the Morlet wavelet for each separate band from 2 to 107 Hz (8 scales/octave).

3.2.5 Electrophysiological relationships with reaction time

To assess the single-trial correlation between reaction time and normalized power changes, the theta power of each trial was averaged across the time period 300 ms prior to the response and correlated with the reaction time in that trial. The resulting correlation coefficients (positive and negative coefficients derived using Spearman's Correlation) were then averaged across the 26 STN sides and a two-tailed, one sample t-test was performed to determine if the mean correlation was significantly different from zero across subjects. In order to determine the presence of a time on task confound, the slope of the RT-theta relationship (θ_{RT}) during the low conflict trials of each subject was used to adjust the theta power of that subject (adjusted_low_conflict_theta= low_conflict_theta + θ_{RT} *($RT_{slow\text{-}high_conflict}$-$RT_{_low_conflict}$)) (Carp et al., 2010). As in the correlation analysis, the 300 ms preceding the response were used to assess significant differences.

The correlation between reaction times and cue-aligned ITPC was done by finding the maximum ITPC value (calculated across all correct trials regardless of congruency) from stimulus onset until 0.5 second after the cue for each channel, averaging across all 6 channels for each subject, and then correlating these values with average reaction time across subjects using Spearman's Correlation. For the response-aligned correlation, the same procedure was done using the maximum ITPC value over a time period from 500ms before to 200ms after the response.

3.3 Results

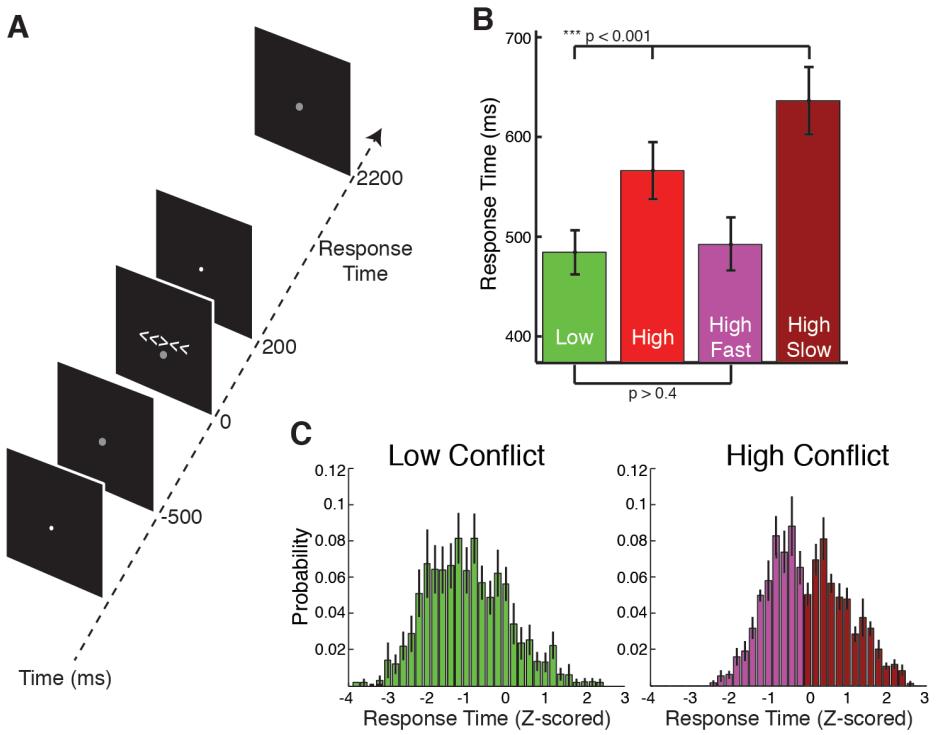


Figure 3.1. Flanker task.

(A) Task. Each trial began with a warning cue 500 ms before arrows were shown. Arrows were shown for 200 ms, and subjects had 2.2 seconds to respond before the next warning cue. A 2:1 ratio of high to low conflict trials was used. **(B)** Behavioural responses across subjects. There was a significant difference between mean low conflict (green bar with SEM) and high conflict (red) reaction times. When the high conflict trials were median split into the fastest half (light red) and the slowest half (dark red), there was no difference in reaction time between the fastest high conflict trials and the low conflict trials. The slowest high conflict trials, however, had reaction times that were significantly slower than the low conflict trials. **(C)** Average reaction time histograms normalized to each subject's mean high conflict trial reaction time. High conflict trial histogram reveals two peaks. Colour denotes whether trials in a given bin were put in the fast-high conflict (light red) or slow-high conflict (dark red) groups for further analysis. Reproduced with permission from (Zavala et al., 2013).

3.3.1 Power changes during correct trials

A schematic of the task performed by the patients is shown in Figure 3.1A. Subjects indicated the direction of a middle arrow that was flanked by either low (<<<<<) or high (>><>>) conflict arrows. In accordance with the literature (Eriksen and Eriksen, 1974), the mean reaction time for correct responses was significantly faster for low conflict trials relative to the high conflict trials (Figure 3.1B; 484 vs. 565 ms, $p<0.0001$, paired t-test). Cue-aligned STN LFP signals in both low and high conflict trials showed an increase in theta power, a decrease in beta power followed by a post response beta rebound, and an increase in gamma power (Figure 3.2A,B). Power differences between low and high conflict trials revealed higher pre-response, conflict related activity

focused in the theta (4-8 Hz) band (Figure 3.2C,D; $p<0.05$). The higher pre-response theta power was also evident in the response-aligned data (Figure 3.2H; $p<0.05$). The effect, however, was weaker than with cue-aligning, as evidenced by the lack of significance when the power difference was corrected for multiple comparisons across the entire spectrogram (Figure 3.2G).

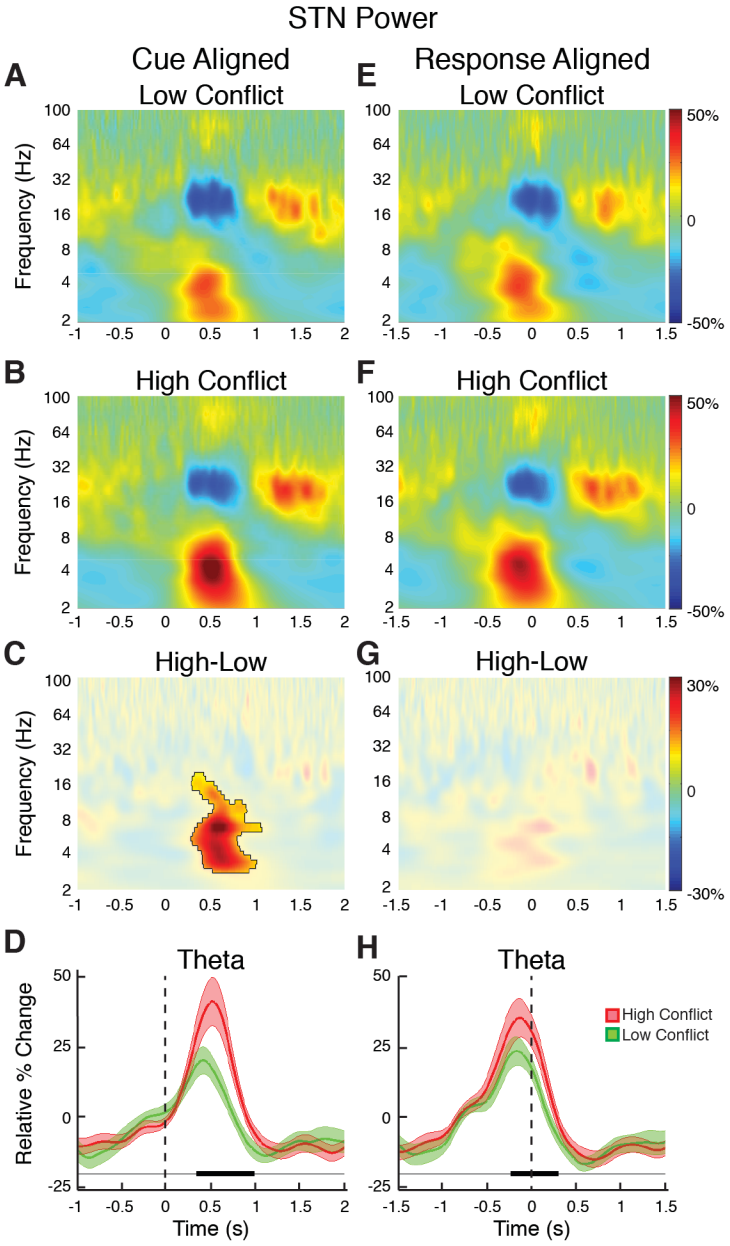


Figure 3.2. Effects of conflict on LFP.
(A-D) Imperative cue-aligned (t = 0) averages (n=26 STNs) of induced spectral change. Both low **(A)** and high **(B)** conflict trials showed an increase in cue-aligned theta power, a decrease in beta power followed by a post response rebound, and an increase in gamma power. **(C)** Difference between trials showed a significant theta band difference (unmasked area, $p<0.05$ after correcting for multiple comparisons, permutation test). **(D)** Cue-aligned theta (4-8Hz)

band average time series (mean ± SEM) for low (green) and high (red) conflict trials. Significant difference between the two conditions is marked by black bar ($p<0.05$ after correcting for multiple comparisons, permutation test). **(E-H)** Same as A-D but aligned to the response. Theta difference is weaker and only significant in the theta band average time series (H). Note that here and in ensuing time-frequency plots frequency is given on a log axis. Reproduced with permission from (Zavala et al., 2013).

We used a 2:1 ratio of high to low conflict trials, and therefore were able to divide the trials into three groups containing a roughly equal number of trials: low conflict trials, the fastest half of the high conflict trials, and the slowest half of the high conflict trials (Figure 3.1B,C). A histogram of normalized reaction times for the high conflict trials provided further support for this procedure by suggesting a bimodal distribution of high conflict reaction times (Figure 3.1C). Figure 3.3 contrasts the STN LFP time-frequency profiles between trial types. When the low conflict trials were compared to the fast-high conflict trials there was no significant difference in reaction time (Figure 3.1B) or induced power (Figure 3.3A,D,E,H). On the other hand, when the slow-high conflict trials were compared to the low conflict trials, slower reaction times were observed (Figure 3.1), and induced theta power was higher in the slow-high conflict trials (Figure 3.3B,D,F,H). The theta power differences were also observed when the slow-high conflict trials were compared to the fast-high conflict trials (Figure 3.3C,D,G,H). Aligning the data to the response revealed that the higher theta power observed during the slow-high conflict trials was mostly confined to the pre-response period (Figure 3.3F-H).

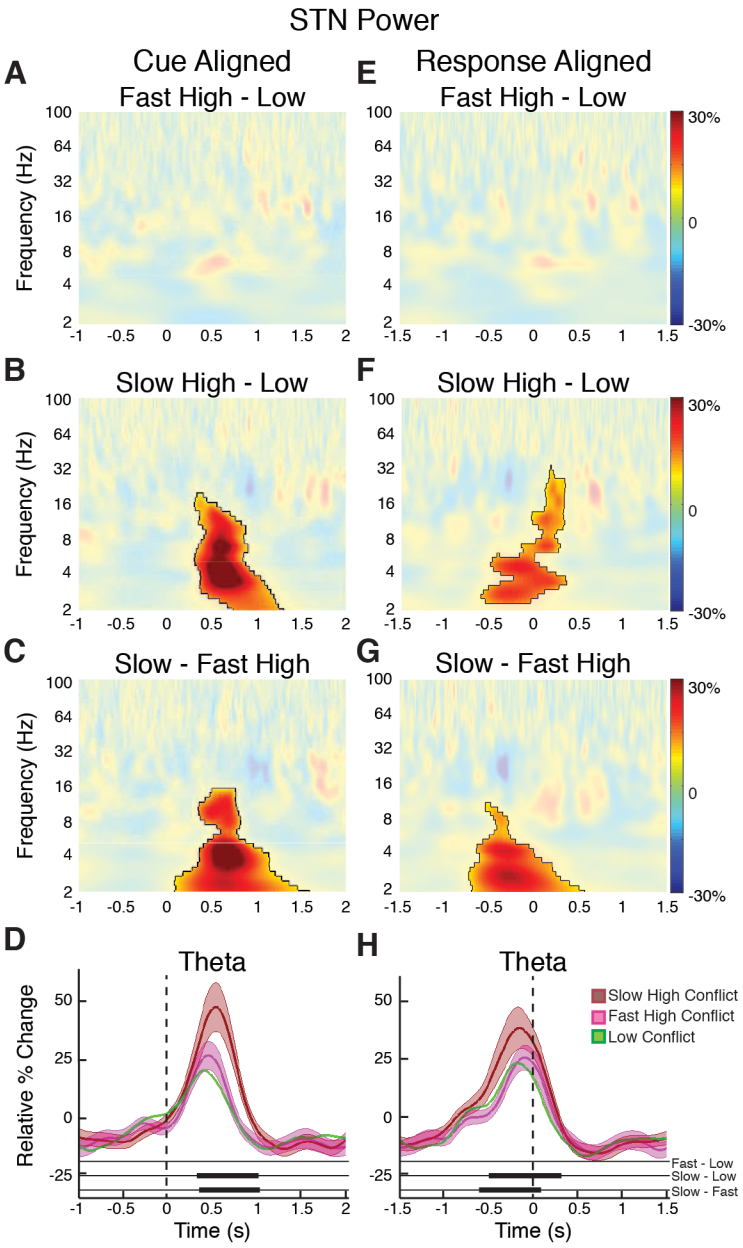


Figure 3.3. Differences between slow and fast trials following high conflict cues.
(A-D) Imperative cue-aligned (t = 0) averages (n=26 STNs) of induced spectral change. **(A)** No power differences between the fast-high conflict trials and the low conflict trials (masked at $p<0.05$ significance level after correcting for multiple comparisons, permutation test). Slow-high conflict trials showed higher cue-aligned theta power than low conflict trials **(B)** and fast-high conflict trials **(C)**. **(D)** Theta (4-8 Hz) band average time series for slow-high conflict (dark red), fast-high conflict (light red) and low conflict (green) trials. Note mean ± SEM are shown except for low conflict trials (where ± SEM were shown in **Figure 3.2**). Significant difference between trial types is marked by horizontal bars ($p<0.05$ corrected for multiple comparisons, permutation test). **(E-H)** Same as A-D but aligned to the response. Fast-high conflict trials showed no significant difference from the low conflict trials. Pre-response theta power was higher in slow-high conflict trials. Reproduced with permission from (Zavala et al., 2013).

Figure 3.3 suggests that the differences between low and high conflict trials shown in Figure 3.2 were due to the slow-high conflict trials. In order to test if there was a relationship between theta

power and reaction time at the single trial level, we correlated the reaction time of each trial to the mean theta power in the 300ms preceding the response. This revealed a significant, albeit low positive correlation for both the low conflict (mean $r=0.17\pm0.03$, $p<0.001$, single sample t-test) and the high conflict trials (mean $r=0.08\pm0.04$, $p<0.05$, single sample t-test). The strength of the correlation was not significantly different between the low conflict and the high conflict trials ($p=0.07$, paired t-test).

Though the correlations we observed are small, they raise the possibility of a time-on-task effect (Carp et al., 2010). The interpretation of this effect is a subject of debate (Nachev, 2011; Yeung et al., 2011; Cohen and Nigbur, 2013; Scherbaum and Dshemuchadse, 2013), but at the very least it serves to highlight a potential ambiguity in interpretation. Are the higher theta power levels due to a longer reaction time that allows theta to continue increasing for longer during conflict, or does the higher theta power reflect the level of conflict, which is indexed by reaction time (Yeung et al., 2004)? Indeed when the theta power observed in the low conflict trials was adjusted to account for the longer reaction times, the high conflict trials were no longer significantly different (high conflict vs. adjusted low conflict= $0.41\pm.09$ vs. $0.42\pm.10$, normalized theta power change ± SEM, $p>0.05$, paired t-test, Carp et al., 2010, see section 3.2.5). Though this suggests that a time on task confound may underlie some of the higher theta power levels observed during conflict, our subsequent work suggests that theta power is higher during conflict even in tasks that do not involve a time on task confound (see Chapter 4 below).

3.3.2 Theta phase inter-trial phase locking during correct trials

Several flanker paradigm studies have demonstrated that cortical theta frequency oscillations demonstrate phase reorganization during the task (Cavanagh et al., 2009; Cohen and Cavanagh, 2011). Accordingly, we tested if a similar phenomenon occurred in the STN. Figure 3.4A-B illustrates the average cue-aligned wavelet ITCP (see methods section 2.2.3 "Oscillatory Inter-trial phase consistency (ITPC)") for the fast-high conflict (Figure 3.4A) and the slow-high

conflict trials (Figure 3.4B). Subtracting the cue-locked ITPC matrices across subjects revealed a significantly lower inter-trial theta phase locking during the slow-high conflict trials (Figure 3.4C). In order to explore this phenomenon further, we band passed filtered the data between 4-8 Hz and compared the ITPC across the three conditions (low conflict, fast-high conflict, and slow-high conflict) using the Hilbert transform (Figure 3.4D). All trial types demonstrated a small ITPC increase following the warning cue (t=-500 ms) and a second, larger increase upon presentation of the arrows, peaking 140 to 200 ms after the arrows were shown. Interestingly, slow-high conflict trials demonstrated a significantly lower ITPC upon arrow presentation than both the low conflict trials and the fast-high conflict trials (Figure 3.4A-D). These differences were significant when the point-wise ITPC values were compared (Figure 3.4D) and when the maximum ITPC value between 0-500 ms was compared ($p<0.01$ for slow versus fast-high conflict and slow-high conflict versus low conflict; $p>.2$ for fast-high conflict vs. low conflict, paired t-tests). This suggests that during the slow-high conflict trials, the stimulus triggered decrease in theta phase variability was impaired.

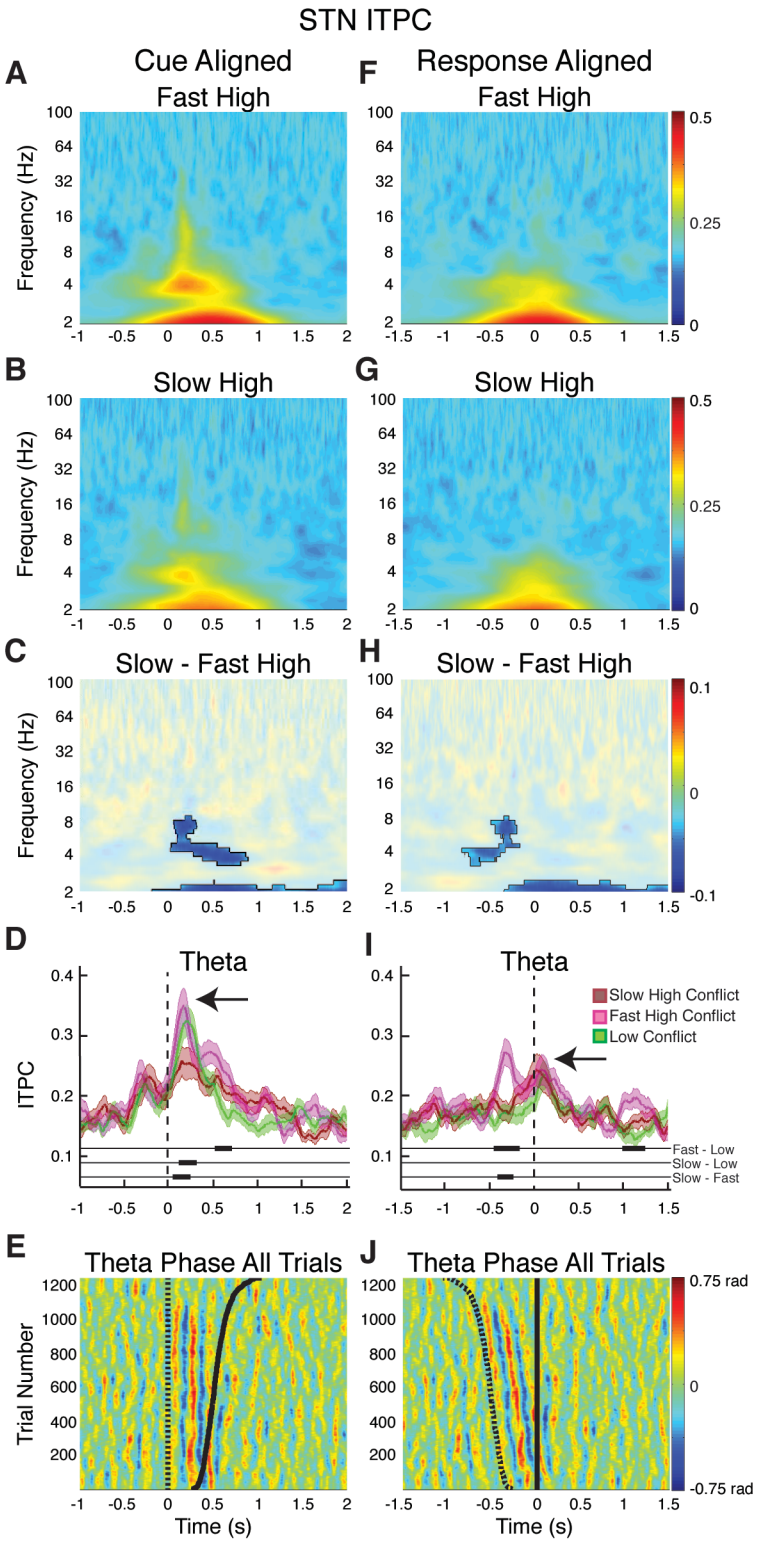


Figure 3.4. Cue-locked theta phase realignment is disrupted in slow-high conflict trials.
(A-E) Imperative cue-aligned (t = 0) averages (n=26 STNs) of phase locking across trials. **(A)** Mean wavelet ITPC for all fast-high conflict trials averaged across all 26 STNs. **(B)** Same as A for slow-high conflict trials. **(C)** There is a significant reduction in cue-locked theta ITPC when the slow-high conflict trials are compared to fast-high conflict trials (masked at $p<0.05$ significance level after correcting for multiple comparisons, permutation test). **(D)** Theta band

filtered Hilbert ITPC also showed impaired cue-locked phase alignment in the slow-high conflict trials. Slow-high conflict (dark red), fast-high conflict (light red) and low conflict (green) mean ITPC time series are shown ± SEM. Significant differences are denoted by horizontal bars ($p<0.05$ after correcting for multiple comparisons, permutation test). **(E)** Theta band-pass filtered Hilbert phase for all trials in all subjects sorted by reaction time (smoothed across each 100 trials and 40 ms). Prominent phase alignment is stimulus-locked and independent of reaction time. Dotted black trace indicates stimulus onset, and solid black trace indicates trial reaction time. **(F-J)** Same as A-E but aligned to response. Comparison of (D) and (I) suggests that there are two phase alignment periods: just after but time-locked to stimulus onset (arrowed large peak in D) and peri-response (arrowed second peak centered on response at t = 0 in E). Both peaks can be seen in the response-aligned Hilbert ITPC of the fast-high conflict trials, which showed the least variable reaction time (I). Cue-locked phase alignment (D) is greater than response-aligned phase alignment (I) for all trial types except slow-high conflict trials, which show impaired cue-locked phase alignment. Reproduced with permission from (Zavala et al., 2013)

Several lines of evidence strongly suggest that the impairment of phase realignment in the slow-high conflict trials was not due to their longer and potentially more variable reaction times. First, the low conflict trials showed a higher peak ITPC despite having a much more variable reaction time than the slow-high conflict trials (STD mean± SEM: 81±6 vs. 57±6; $p<1*10^{-7}$, see Figure 3.1C for a comparison of the distributions). Secondly, when the low conflict trials were median split and analysed in the same way as the median split high conflict trials, slow and fast low conflict trials showed no difference in their cue-locked theta ITPC increase (neither point-wise comparisons nor comparisons of the maximum ITPC value showed any differences, data not shown). This suggests that the difference between the fast and slow-high conflict trials was not simply due to reaction time differences. Finally, the phase alignment was locked to the cue and peaked approximately 300ms prior to the average response, which occurred between 485-640 ms after the cue depending on trial congruency (Figure 3.4). This cue locking was particularly well seen when the theta phase from all correct trials recorded in all patients was aligned to the cue (Figure 3.4E, sorted by reaction time, channel with the clearest cue-locked phase alignment was used for each patient). These showed a prominent increase in phase synchrony locked to the presentation of the arrows at time t=0.

In order to verify that the phase alignment just after the imperative cue (arrows) was indeed locked to the cue rather than the response, we repeated the same analysis with the data locked to the response. This revealed that there was another, weaker period of phase realignment that occurred when subjects made their response. The response-locked increase in phase locking

peaked between 20-100 ms after the response was made, which was about 400 ms after the peak in the cue-aligned ITPC increase (Figure 3.4F,G,I). This suggests that the cue-aligned ITPC increase and the response-aligned ITPC increase were two separate events. This is supported by the fact that two peaks can be seen in the response-aligned fast-high conflict trials (Figure 3.4I). Both peaks are most likely visible only in the fast-high conflict trials because these trials had the least variable reaction times (fast-high conflict vs. slow high conflict mean STD±SEM: 39±2 vs. 81±6 ms, $p<1*10^{-8}$; fast-high conflict vs. slow-high conflict mean STD±SEM: 39±2 vs. 57±6 ms, $p<0.01$). In fact, the response triggered increase in ITPC for the fast-high conflict trials may also potentially be responsible for the second, small relative increase in the cue-aligned fast-high conflict trial data of Figure 3.4D (blue trace at around t=500 ms), where it is superimposed on the falling end of the earlier cue-locked ITPC peak.

Further evidence that the cue-locked increase in theta ITPC is not due to 'smearing' from a response-locked theta ITPC increase stems from the fact that the former is larger than the latter (compare black arrows in Figure 3.4D,I). When the theta phase from all correct trials recorded in all patients was aligned to the response (Figure 3.4J, sorted by reaction time), a diagonal trend parallel to the onset of each trial (dotted black trace) could be seen. This suggests that the dominant theta phase alignment effect was locked to the cue and not to the response. In line with this, the cue-aligned peak was significantly higher than the response-aligned peak when the ITPC was calculated across all trials (regardless of arrow congruency) and across all low conflict trials, all high conflict trials and the fast-high conflict trials ($p<0.05$ for all comparisons, paired t-test). The only trials that did not show a significant difference in the peak ITPC between the cue and the response-aligned data were the slow-high conflict trials. This was due to the disruption of the cue-aligned phase-locking during the slow-high conflict trials as mentioned above.

As the finding that the theta phase realigned during the stimulus onset was quite novel, we dedicated a significant amount of effort proving that it was indeed separate from any events

related to movement. Our subsequent work (Chapter 5 and Chapter 6) has further validated this finding and attempted to explore its physiological purpose.

Unlike the cue-aligned increase in phase-locking, the maximum value of the response-aligned increase (between -200 ms to 200ms) showed no difference between trial types ($p>0.05$, paired t-test). This together with the timing of the two peaks suggests that only the cue-aligned ITPC increase may play a functional role in determining reaction time. This is supported by the finding that when the ITPC was computed across all correct trials (regardless of congruency), the maximum cue-aligned ITPC correlated with reaction time across subjects (Figure 3.5A, $r=-0.58$, $p<0.05$). When the data were aligned to the response, on the other hand, there was no significant correlation between peak ITPC and reaction time (Figure 3.5B, $r=-0.11$, $p>0.05$).

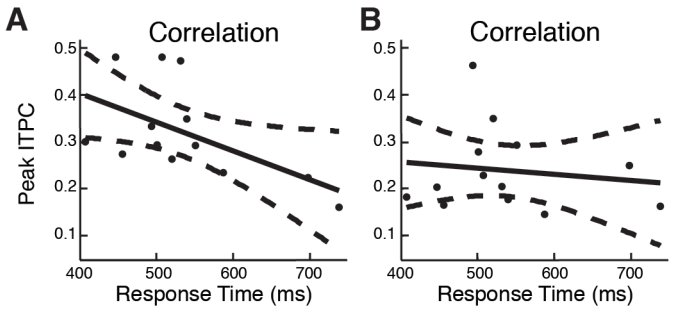


Figure 3.5. Cue-locked phase alignment correlates with reaction time.
(A) Cue-locked phase alignment (calculated across all correct trials) correlated with subject reaction time ($r=-0.58$, $p<0.05$). **(B)** Response-locked phase alignment did not correlate with subject reaction time ($r=-0.11$). Correlations show results of linear regression and corresponding 95% confidence limits. Reproduced with permission from (Zavala et al., 2013)

3.3.3 Power changes in error trials

As a final analysis of this data set, we contrasted the spectral changes between the correct high conflict and the incorrect high conflict trials (16.4 ± 2.8 %; only 4.9 ± 1.6 % of low conflict trials were incorrect, -too few to perform such an analysis). Incorrect high conflict trials had faster reaction times than correct slow-high conflict trials but not correct fast-high conflict trials (Figure 3.6A). As with the correct trials, the incorrect trials demonstrated an increase in theta power, a decrease in beta power, and an increase in gamma power (Figure 3.6B). When the incorrect-high

conflict trials were compared to the correct-high conflict trials, a higher delta power (2-4 Hz) was seen in both the cue and response-aligned data (Figure 3.6B-E) and a significant increase in 10-20 Hz power after the response in the response-aligned data (Figure 3.6D,E).

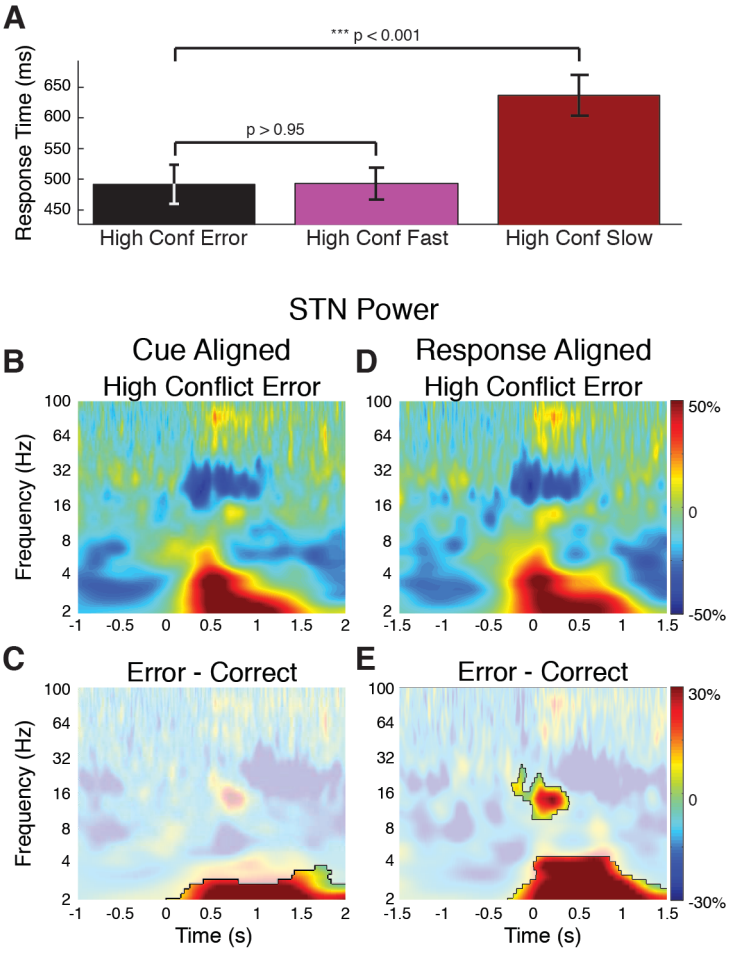


Figure 3.6. LFP in error trials across all subjects.
(A) Incorrect high conflict trials had a group mean reaction time (±SEM) that was faster than slow-high conflict trials, but no different than fast-high conflict trials. **(B)** Average (n=26 STNs) induced LFP power changes during incorrect high conflict trials aligned to imperative cue. **(C)** Differences between incorrect and correct high conflict trials across subjects (unmasked area, $p<0.05$ after correcting for multiple comparisons, permutation test). **(D-E)** Same as B-C but aligned to the response. Increases in post response delta (2-4Hz) and low beta (10-20 Hz) power can be seen. Reproduced with permission from (Zavala et al., 2013)

3.3.4 Spatial focality of recordings

Here we present the results for the average across all contract pairs of the deep brain stimulation electrode targeting the STN on different sides across patients. LFP recordings from the subthalamic region were highly focal, as indexed by steep percentage drops in power when

comparing the contact pair with the highest absolute power to the mean power recorded by the two remaining contact pairs on each electrode. Relative to the best contact pair, the relative mean theta power for the two remaining channels dropped to 48.4 ± 4.1% (mean±SEM). When the low versus high conflict trial comparisons were repeated using only the channel with the highest beta or theta power, the same relative increase in theta power and decrease in theta ITPC were observed during the slow-high conflict trials (data not shown). Though the variability of contact locations across each STN did not allow us to find any significant topographical segregation of the reported effects on theta power and phase, we found that on average the contact with the highest beta power was less deep than the contact with the highest theta power by 0.31 times the inter-contact distance (or 0.62 mm if we take this distance as that between the centre of each contact). Although this difference is compatible with the hypothesis that theta activity is more focused in associative and limbic STN and beta activity more focused in dorsolateral motor STN, the difference was not statistically significant ($p=0.151$, Wilcoxon rank sum test).

3.4 Discussion

Our results suggest that STN LFP activity tends to present one of three spectral variants depending on Flanker trial type. Correct low conflict and fast-high conflict trials demonstrated similar theta phase realignment when the arrows were shown, followed by a comparable theta power increase. The similarity in LFP response between the two trial types was associated with similar reaction times and may imply that subjects were able to appropriately ignore the distracting flankers during correct fast-high conflict trials. The slowest correct high conflict trials, where we presume that the distracting arrows were attended to resulting in conflict, were associated with a disruption of cue-aligned theta phase reset, a higher theta power increase, and longer reaction times. The theta power observation is consistent with other studies that have also reported higher STN theta power during high conflict trials (Cavanagh et al., 2011; Fumagalli et al., 2011; Brittain et al., 2012) but the lower theta phase reset is less consistent with some of our

own subsequent work (see Chapter 5 and Chapter 6). The final variant of STN spectral activity occurred during incorrect responses, which showed higher delta (2-4Hz) and alpha/low beta (10-20 Hz) power, but only during or after the response. Error related delta power increases have previously been observed over mPFC (Yordanova et al., 2004; Cavanagh et al., 2012; Cohen and van Gaal, 2014).

3.4.1 Preliminary hypotheses for how STN spectral features help determine response timing under conditions of conflict

We have identified two temporally separate events in the STN reflected in the LFP. The first is a theta band phase re-organization that follows the imperative cue at short latency and is bigger in low conflict and fast-high conflict trials than in slow-high conflict trials. This feature is phase locked to the imperative cue and not the response, as demonstrated by the twofold decrease in average magnitude when trials are re-aligned to response as opposed to stimulus onset (also see chapters Chapter 5 and Chapter 6 for further validation that the theta phase alignment is indeed locked to the stimulus and not the response). These observations raise the possibility that the theta band phase reorganization reflects early cortical and subcortical activity related to the cue that nevertheless then impacts the speed with which the relevant behaviour is processed. Recently, Cohen and Cavanagh (2011) have shown that mPFC theta phase in the flanker task correlates with reaction time up to 500 ms before a response occurs. In other words, some mPFC theta phase values predicted long reaction times and others predicted short reaction times. Here, we have demonstrated that the STN shows theta phase realignment when an imperative cue is given and that the degree of alignment across trials negatively correlates with reaction time across subjects. One way theta phase resetting could affect function is through promoting or disrupting optimum phase alignment with relevant afferent and efferent targets (Fries, 2005). For example, studies have shown that the degree of theta phase coupling between the mPFC and inferior frontal cortex following errors correlates with post-error slowing on the following trial (Cavanagh et al.,

2009; Cohen and Cavanagh, 2011). In the context of communication through coherence (Fries, 2005), the disruption of STN phase alignment during slow-high conflict trials may reveal a deviation from the ideal phase for rapid response execution. Though we are unsure how conflicting stimuli could attenuate theta phase re-organization, one possibility is that during the trials in which conflict has the strongest effects other brain structures that respond to conflict become synchronized with the STN and therefor "hijack" the theta oscillations. We further explored this possibility in subsequent studies (see Chapter 4 and Chapter 6).

The second event is an increase in STN theta band power. It follows the cue-locked theta band phase re-organization and peaks just before the response. Unlike the theta phase re-organization, its magnitude is not significantly reduced when trials are re-aligned to response onset, as opposed to cue onset. Together these observations suggest that theta power is at least partially linked to the response itself, possibly helping to determine its timing. Also in contrast to the theta phase re-organization, the theta power increases were higher in the slowest trials of the high conflict condition.

There are several valid interpretations of why theta power, both in the STN and in the mPFC (Cohen and Cavanagh, 2011; Cavanagh et al., 2012; Cohen and van Gaal, 2014; Oehrn et al., 2014), could be higher during conflict. The simplest explanation that has been proposed is that theta band activity simply reflects time spent on a task (Carp et al., 2010; Scherbaum and Dshemuchadse, 2013). If theta power correlates with reaction time and high conflict trials have longer reaction times, theta power should be higher during conflict. However, several authors have argued against a simple time on task interpretation (Yeung et al., 2011; Cohen and Nigbur, 2013) by citing studies showing that theta-RT correlations are stronger during high conflict (Cohen and Cavanagh, 2011) trials relative to low conflict trials and that error trials, which typically have short reaction times, also have high theta power increases (Cavanagh et al., 2009;

van Driel et al., 2012). Furthermore, our subsequent work described below in Chapter 4 suggests that STN theta power increases reflect more than just time on task (Zavala et al., 2014).

If theta is not simply an indicator of time on task, then what could it be? Though we run the risk of overly simplifying a complex network, one way of trying to address this issue is to ask whether STN and mPFC theta power increases reflect anti- or pro-kinetic activity. Much of the literature would suggest that theta activity is anti-kinetic. First, the traditional firing rate model (Albin et al., 1989; DeLong, 1990) would predict that increased activity in the STN would excite the GPi and thus inhibit the thalamus and subsequently inhibit movement (see Figure 1.1). Furthermore, in agreement with the correlations we report here, higher mPFC theta power correlates with longer reaction times and higher evidence thresholds (Cavanagh et al., 2009, 2011; Cohen and Cavanagh, 2011), again suggesting that it is inhibitory. It could be the case that during high conflict trials, the mPFC detects conflict and subsequently activates the STN via theta band oscillations to inhibit movement. It is important to stress, however, that the low conflict trials also showed an increase in theta band power. Viewing STN theta activity as strictly anti-kinetic would then suggest that the low conflict flanker trials also elicit conflict (perhaps due to the possible presence of conflict on any trial). Alternatively, the theta power increases we observe may correspond to activity responsible for inhibiting all incorrect actions, something that needs to be done on all trials regardless of the congruency of the flanking arrows. According to this interpretation, the higher increases in theta power that occur late during high conflict trials may reflect a greater drive needed to inhibit the incorrect response due to a stronger activation of that response by the flanking arrows.

However, despite the location of the STN in the context of the rest of the basal ganglia (see Figure 1.1 and section 1.1.2 “The Basal Ganglia Rate Model: The direct and indirect pathways”), there is no reason to assume that theta activity reflects exclusively anti-kinetic inputs, especially considering that the STN receives direct inputs from pre-motor areas that also activate the

striatum (Nambu et al., 1996). Indeed, some models have predicted that the STN inhibits movement during conflict by integrating all of the "pro-kinetic", pre-motor cortex activity in favour of all possible actions during a given task and activating the GPi proportionately (Frank, 2006; Bogacz and Gurney, 2007). When viewed in the light of this hypothesis, one possible interpretation of our results could be that STN theta band activity serves as a marker of the basal ganglia's activation by pre-motor cortical neurons. This may help explain why theta power continues to rise until the response is executed; anti-kinetic activity would be more likely to drop before the response to allow for movement. From this perspective, the higher theta activity observed during conflict would be due to the increased input to the basal ganglia for both the correct and the incorrect response. This interpretation may help explain why Cavanagh et al found that STN DBS didn't just remove the correlation between mPFC theta power and evidence threshold, it reversed it (Cavanagh et al., 2011). Higher mPFC theta power was therefore associated with faster reaction times during DBS.

In conclusion, our first study showed that theta band phases and amplitudes are altered in the STN during the flanker task, and that these changes were further modulated by conflict. The physiological role of these changes however was left open, with several different possible interpretations. Thus subsequent chapters focus on trying to confirm or rule out some of these interpretations.

Chapter 4. Midline frontal cortex low-frequency activity drives subthalamic nucleus oscillations during conflict

4.1 Introduction

As our first study (Chapter 3) suggested, conflict related activity in the STN showed many similarities to that observed over the mPFC. STN theta power was higher during conflict, and the theta power level correlated with reaction time as it does in the mPFC (Cohen and Cavanagh, 2011; Cavanagh et al., 2012; Oehrn et al., 2014). These findings, together with the fact that STN DBS reverses the relationship between conflict-related mPFC theta power and reaction time, suggest that mPFC theta activity drives similar activity in the STN, which then adjusts the threshold of evidence necessary for a response to be made during conflict. In line with this hypothesis, both neuroanatomical and imaging studies have demonstrated connectivity between the mPFC and the STN (Smith et al., 1998; Nambu et al., 2002; Aron et al., 2007; Haynes and Haber, 2013). However, hitherto there has been no proof of a dynamic theta drive from the mPFC to the STN during high conflict scenarios. In order to test the hypothesis that the two structures synchronize their activity in a frequency specific manner, we simultaneously recorded STN local field potential (LFP) activity and midline electroencephalographic activity (EEG) while patients performed a sensory integration task involving conflict. In order to address some of the more nuanced questions related to the interactions between STN theta power increase and time on task (see preceding Chapter 3), we designed a novel task that separated conflict and motor responses in time.

4.2 Materials and Methods

4.2.1 Subjects and Task

Thirteen subjects (seven males, mean disease duration 10 years, mean age 56 years, range 42– 66 years) underwent bilateral implantation of DBS electrodes into the STN according to the

procedures described in the methods section 2.1.2 “Post-operative recordings in the UK”. Patients participated in a decision-making task while on their regular medication three to six days after electrode implantation. All subjects gave their written informed consent to take part in the study, which was approved by the appropriate local ethical committees. Clinical details of the patients are provided in Table 4.1. The mean percentage improvement in the motor section of the Unified Parkinson’s Disease Rating Scale (UPDRS) on treatment with levodopa was 63 ± 4 % across subjects ($p < 0.001$, paired t-test). The improvement in all subjects’ scores with medication indicated a good responsiveness to levodopa in our study participants. Only cases 10 and 13 had been diagnosed as having an impulse control disorder, however, neither showed behaviour that deviated significantly from the other subjects. The mean response time for the two patients with impulse control disorders was 3.182 and 4.594 s, which fell well within the range of the other 11 patients (1.828-4.817 s).

Table 4.1. Clinical details from dots experiment.

Case	Age	Disease Duration	UPDRS Off (III)	UPDRS On (III)	First Symptom	Reasons for Surgery	Daily Medication (mg/d)
1	49	10	42	6	Tremor	Tremor	Levodopa 300; Trihexyphenidyl 2
2	50	9	58	23	Should Stiffness	Motor fluctuations, Tremor	Levodopa 1050
3	50	4	N/A	N/A	Right arm tremor	Tremor	Levodopa 400; Rotigotine 16; Entacapone 600
4	66	16	32	13	Loss of dexterity	Bradykinesia	Levodopa 600; Amantadine 200; Ropinirole 24; Rasagiline 1
5	51	7	58	13	Loss sense of smell	Tremor, Gait difficulties	Levodopa 1300
6	64	12	70	20	Tremor	Dyskinesias	Levodopa 1200; Apomorphine 7mg/hr
7	47	14	34	11	Left arm bradykinesia	Dyskinesias, Motor fluctuations	Levodopa 350; Pramipexole 1.05; Amantadine 300
8	66	14	63	24	Shoulder pain & stiffness	Motor fluctuations	Levodopa 650; Pergolide 9
9	57	7	43	17	Left hand tremor	Tremor	Levodopa 600; Rotigotine 8; Ropinirole 8; Rasagiline 1; Entacapone 1200
10	57	6	21	7	Left side bradykinesia	Dyskinesias, Motor fluctuations	Levodopa 750; Entacapone 1000
11	61	4	37	15	Left side	Tremor	Amantadine 200;

					tremor		Levodopa 750; Entacapone 1000
12	65	15	51	21	Left hand tremor	Freezing	Amantadine 200; Levodopa 400; Ropinirole 12
13	42	9	60	42	Loss of dexterity	Bradykinesia, dystonia, freezing	Amantadine 400; Levodopa 600

A schematic of the task is shown in Figure 4.1A. Subjects were asked to stare at a target circle containing 200 randomly moving white dots presented on a black background. The task was generated using the program Psychopy (Peirce, 2007) and presented on a 33 cm Macintosh laptop (60 Hz screen refresh rate). The target circle was 14 cm in diameter and the white dots were 10 pixels large (visual angle ≈ 0.25 °). Prior to the beginning of the experiment, the subjects were asked to indicate the overall direction the dots were moving whenever they noticed that the dots were coherently moving either left or right. Left direction choices were indicated with a left fore-finger press of the "z" key on the keyboard, and right direction choices were indicated with a right fore-finger press of the "/" key. Each randomly moving dot moved in a straight line at a rate of 0.14 mm per frame for 20 frames (333 ms) before moving to another part of the target circle and moving in a new direction chosen pseudo-randomly between -180 and 180 degrees. This was done to prevent the subjects from focusing on any individual dot and using it to identify coherent motion of the dots.

Two types of trials were randomly shown throughout the experiment. Half of the trials (40 trials) were what we will refer to as "low conflict trials". At the beginning of all low conflict trials, some of the dots started moving horizontally to either the right (0 degrees) or the left (180 degrees) of the screen. The number of dots moving coherently in one direction increased linearly from 0% to 50% at a rate of 0.002% per frame. By 4.16 s, the maximum number of dots moving coherently had been reached (100 dots, 50% of all 200 dots). The second half of the trials (40 trials) were what we will refer to as "high conflict trials". The high conflict trials were identical to the low conflict trials with the exception that at the same time that some of the dots started moving

coherently in one direction, a second group of dots started moving coherently in the opposite direction. Though the percentage of dots moving coherently in either direction increased at the same rate, the maximum number of dots moving in the "conflicting" direction was capped at 20 (10% of all 200 dots, Figure 4.1A bottom panel). By 0.833 s, the maximum number of dots moving coherently in the "conflict" direction had been reached. Regardless of trial type, subjects were instructed to indicate the direction of the majority of the dots. If the subject answered incorrectly or took longer than 14 s to indicate a response, the word "incorrect" was briefly displayed in the center of the screen for 0.75 s. Incorrect trials were excluded from all analysis. Subjects were allowed to practice the task as long as they wished before the electrophysiological recordings were made. The practice sessions were generally quite short (<10 trials) as the task was designed to be as simple as possible.

It is important to note here that at no point were the subjects given a "cue" indicating that a trial had begun. In other words, there was never a defined stimulus onset and all increases in coherent dot movement were gradual. After the subject pressed a keyboard key to indicate his or her response for a given trial, all dots immediately began to move in random directions. They continued to do so for a pseudo-randomly chosen time period of no less than 2 s and no more than 4 s before slowly starting to move coherently again for the next trial. Only when the subject noticed that the dots were moving coherently to the left or the right did the subject become aware that the trial had begun. These gradual changes were designed to suppress neuronal activity caused by abrupt stimulus changes and to increase the separation in time of neuronal activity associated with conflict processing from activity associated with movement (O'Connell et al., 2012).

For the comparison of correct low and high conflict trials, all incorrect responses (11 ± 3% of all trials) as well as any trials with reaction time greater than 8 s or shorter than 150 ms were discarded (6%). From the remaining trials, any trial with a response time that was greater than

two standard deviations (SDs) above the mean latency of the other trials was discarded. In order to determine whether there was a significant difference in reaction time between the two trial types on an individual subject level, the difference of the mean high conflict trial reaction time and the mean low conflict trial reaction time was compared to 1000 surrogate differences generated by randomly permuting the low and high conflict trials. In order to determine significance across subjects, the mean difference across subjects between the two trial types was compared to the distribution of 1000 surrogate mean differences generated by randomly permuting the low and high conflict trials of each subject before calculating the mean difference across subjects. In order to determine if there was a significant across-subject difference in error rates between low and high conflict trials, the mean error rate difference across subjects was compared to 1000 surrogate mean error rate differences generated by randomly permuting each subject's error rate for the high and low conflict condition before calculating the mean error rate difference across subjects.

4.2.2 Electrophysiological data recording and analysis of power

STN LFPs were recorded from the DBS electrodes. Simultaneously, continuous scalp electroencephalogram (EEG) was recorded from frontocentral, central and parietal electrodes at the midline (Fz, Cz and Pz, International 10-20 System). More lateral electrodes were prohibited by surgical wounds and dressings in this patient group. All signals were sampled at 2048 Hz, band-pass filtered between 0.5 and 500 Hz, and amplified using a TMSi porti and its respective software (TMS international, Netherlands). Data were analysed using custom written Matlab (Mathworks Inc., USA) scripts. Prior to any analysis, monopolar recordings were down-sampled to 1000 Hz, notch filtered at 50 Hz, and converted to a bipolar montage between contacts (three bipolar channels per STN side and two bipolar channels for the EEG recordings: Fz-Cz, and Pz-Cz). Re-referencing was done so as to limit the effects of volume conduction from distant sources. Any trial with a clear artefact in either the EEG or LFP bipolar traces was discarded

(mean final trial count after removing artefact trials, incorrect response trials, and behavioural outliers = 56 ± 3.8 trials per subject). For one subject (case 2), the entire recording had to be discarded due to artefacts. One subject (case 1) was only implanted in one STN; therefore the total number of STNs studied was 23. In two of the STNs, only one bipolar channel was available because one of the electrode's middle two contacts was not active.

The power and phase of the bipolar channels in each trial were calculated from 4.5 s before to 1.5 s after the response using the Morlet wavelet at 8 scales/octave from 2 to 107 Hz. To assess the differences in induced power between low and high conflict trials, we used the "induced power differences" approach outlined in the methods section 2.2.2 "Oscillatory power analysis" above. Each channel's time-frequency pixels were then normalized to the overall mean power of that frequency across all time points in all trials for that channel. The normalized induced power was then averaged across the three bipolar contacts in each STN, and the difference between the low and high conflict trials was calculated for each STN and for the two EEG bipolar channels. Averaging across all the contact pairs in a given STN electrode was performed so as to avoid selection bias, although this procedure might serve to underestimate spectral changes in the STN as not all contact pairs were necessarily within or bridging this nucleus.

To assess the statistical significance of the mean induced power difference between low and high conflict trials, the same procedure was repeated 1000 times with the low and high conflict labels of each trial randomly assigned during each permutation. The p-value of each time–frequency pixel was found by comparing the actual mean difference to the distribution of the 1000 permutations for that time-frequency pixel. The p-values were then corrected for multiple comparisons using exceedance mass testing (Maris and Oostenveld, 2007), (see methods section 2.3.2 "Permutations and exceedance mass testing" above).

4.2.3 Granger Causality analysis

One of the primary interests of this experiment was to explore cortico-STN electrical coupling during conflict. In order to do so, we utilized multivariate autoregressive (MVAR) modelling to estimate the coherence (Gardner, 1992; Pesaran, 2008) and Granger causality (Granger, 1969; Kamiński and Blinowska, 1991; Ding et al., 2006; Seth, 2010) between LFP signals recorded from the STN and EEG signals recorded over the frontal and posterior cortices (see methods section 2.2.6 "Granger Causality analysis").

4.2.4 Granger Causality differences between low and high conflict trials

In order to test for time evolving differences in STN:cortical coupling between low and high conflict trials, the following procedure was done for each STN:EEG pair (STN:Fz-Cz, STN:Pz-Cz). First, the MVAR matrix was used to estimate the coherence between each LFP-EEG bipolar pair as outlined in the methods section 2.2.6 "Granger Causality analysis". The LFP-EEG coherence was found for each of the 3 bipolar contacts for each STN. The coherence value at each time-frequency window for each STN contact:EEG electrode pair was then normalized to the "baseline" coherence value for that frequency in that STN contact :EEG electrode pair. Low and high conflict trials were normalized separately using the same baseline. The "baseline" at each frequency was determined by calculating the mean coherence value (averaged across all 13 time blocks) for a set of trials randomly chosen from all of the correct trials. The number of trials used to generate this baseline for each subject was equal to the number of correct low conflict trials for that subject. Once the LFP:EEG coherence had been normalized for each of the three contacts in each STN, it was averaged across all three STN contacts giving a mean STN:EEG coherence spectrogram. The mean difference across all STNs was then found by subtracting the normalized STN:EEG coherence spectrogram for the low conflict trials from the normalized STN:EEG coherence spectrogram of the high conflict trials and averaging across all 23 STNs. This mean difference was then compared to 1000 permuted differences found by randomly permuting the low and high conflict trials before calculating the MVAR matrices. Aside from the

shuffling of trials, the 1000 permutations were analysed in the exact same way as the real data; they were normalized to the same "baseline", and the model order used to analyse the permutated data was the same as that used to analyse the real data. In order to correct for multiple comparisons across time-frequency points, exceedence mass testing was used (Maris and Oostenveld, 2007), see methods section 2.3.2 "Permutations and exceedance mass testing" above. When testing for a conflict related difference in DTF coupling, the same normalization and permutation analysis used to test for differences in coherence was applied to the DTF values for the STN-to-cortex direction and for the cortex-to-STN direction.

4.2.5 Hilbert coherence analysis

In order to further validate our results, the cortico:STN coherence was also calculated using the continuous time evolving methods outlined by Lachaux et al (2002) and explained in detail in the methods section 2.2.5 "Inter-site Lachaux coherence". The time evolving coherence signal was then normalized by each channel's "baseline" coherence. The baseline was chosen in the same way as it was chosen for the Granger analysis: calculating the mean coherence value (averaged across the entire time series) for a set of trials randomly chosen from all of the correct trials. The three resulting normalized time series generated for each of the three contacts in each STN were then averaged within each STN before averaging across all 23 STNs. Statistical significance was determined using permutation testing as outlined above. In order to calculate the Hilbert parietal cortico:STN coherence, the same analysis was done using the Pz-Cz bipolar electrode instead of the Fz-Cz electrode.

As a final validation of our results, the same analysis outlined above for the Hilbert coherence was redone using only the phase values i.e. the inter-site phase coherence (see methods section 2.2.4 "Inter-site phase coherence" above). In other words, instead of finding the magnitude of the average cross-spectrum and dividing by the average auto-spectrum at each time point, the

magnitude of the average phase difference between the LFP and the EEG signals was calculated at each time point.

4.2.6 Calculation of power and coherence distributions across STN bipolar contacts

Here we assumed that the bipolar contact pair with the highest beta band power in each electrode was in or nearest to the dorsolateral motor territory of the subthalamic nucleus, as reviewed in (Brown, 2013). Nineteen of the STNs contained at least one bipolar contact that was ventral to the contact with the highest beta. Comparing the changes in power observed over the highest beta contact with the changes in power observed over the remaining ventral contacts revealed that the higher theta power observed during high conflict trials was independent of beta power localization (ANOVA, within-subject repeated measures, conflict x channel: conflict $F=8.02$, $df=1$, $p=0.006$; channel $F=0.02$, $df=1$, $p=0.91$; interaction $F=.01$, $df=1$ $p=0.92$). Similarly, we were unable to show any localization of the higher delta power observed during high conflict trials (ANOVA, within-subject repeated measures, conflict x channel: conflict $F=5.13$, $df=1$, $p=0.03$; channel $F=0.62$, $df=1$, $p=0.43$; interaction $F=.13$, $df=1$ $p=0.71$) or the higher frontal cortico:STN coherence observed during high conflict trials (ANOVA, within-subject repeated measures, conflict x channel: conflict $F=5.7$, $df=1$, $p=0.02$; channel $F=0.05$, $df=1$, $p=0.83$; interaction $F=0.1$, $df=1$ $p=0.77$). Our previous study (see section 3.3.4 “Spatial focality of recordings” above) also failed to reveal any localization of our results to specific subsections of the STN. This, however, does not suggest that the STN is homogenous. Rather, it implies that without high resolution imaging, it is difficult to infer the precise location within the STN using only simple biomarkers such as theta and beta power. The effect is likely compounded by our choice to use bipolar contacts when analysing our data, however, we feel that reducing the effects of volume conduction takes precedent over the inherently indeterminate task of localizing electrodes to subdivisions of a small nucleus like the STN.

4.3 Results

4.3.1 Behavioural results

A schematic of the task performed by the subjects is shown in Figure 4.1A. At the beginning of the task, the subjects were asked to stare at a screen that contained continuously displayed randomly moving dots. Without any explicit cue indicating that a trial had begun, the dots gradually increased their coherent motion on each trial, and the subjects were asked to indicate the direction of overall dot movement whenever it became apparent to them that the dots were moving in one consistent direction. During the low conflict trials, dots gradually increased their coherent horizontal motion to either the left or the right. Dot coherence was increased linearly at a rate of 0.002% per frame (60 Hz refresh rate) until the subject indicated the overall direction of movement or the dots reached a 50 % coherent movement limit (Figure 4.1A, top). During the high conflict trials, the above population of coherent moving dots behaved in the exact same way as they did during low conflict trials. Additionally though, a subpopulation of the remaining dots gradually increased their coherent motion in the opposite direction to the majority, inducing conflict (Figure 4.1A, bottom). Dots moving in the conflicting direction never exceeded those moving in the preferred direction and were capped early in the trial (after 0.83 s) at 10% of the total number of dots on screen.

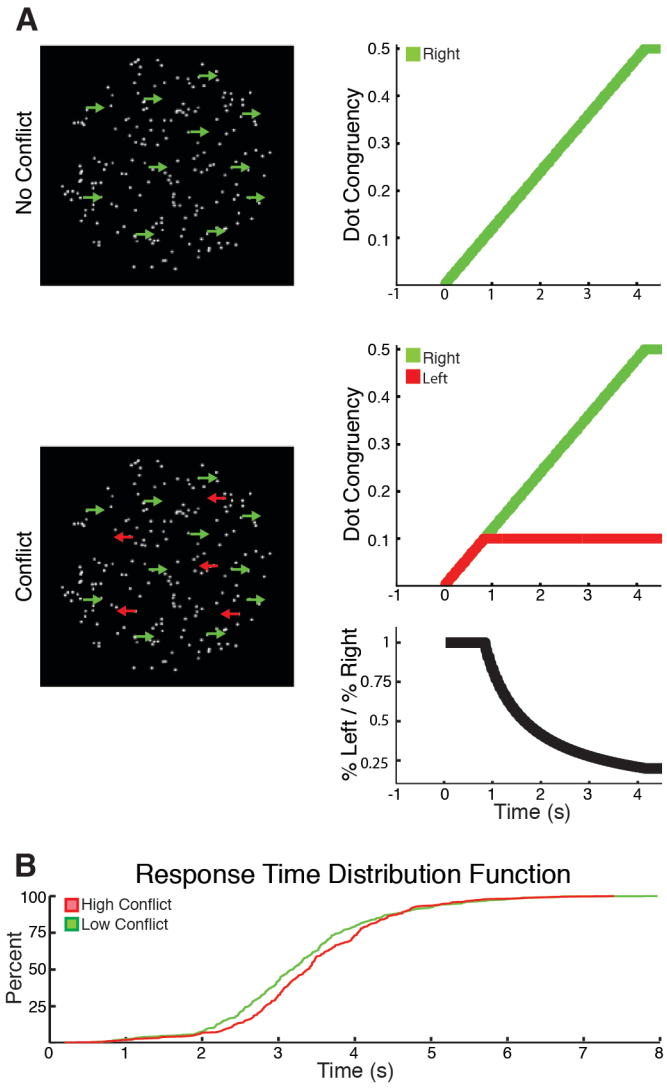


Figure 4.1. Task and behaviour.
(A) Randomly moving dots were displayed continuously on screen. During low conflict trials (A, Top) dots began moving coherently in one horizontal direction until either 50% of the dots moved coherently, the subject indicated the direction he or she perceived the dots were moving in, or the trial timed out after 14 s. During high conflict trials (A, Bottom), up to 10% of the dots gradually moved in the opposite direction to the above population. Conflict was highest during the first 833 ms of the trial and then gradually decreased over time as the number of dots moving in the correct direction exceeded the number of dots moving in the conflicting direction (bottom black trace). High and low conflict trials were randomly interleaved throughout the experiment without the subject knowing what the current trial type was. **(B)** Cumulative distribution function of the response times for all of the low and high conflict trials (collapsed across all subjects) used in the analysis. Reproduced with permission from (Zavala et al., 2014).

On average, subjects required 3.29 ± 0.23 s (SEM) to respond on low conflict trials and 3.43 ± 0.23 s to respond on high conflict trials ($p<0.05$, permutation testing, Figure 4.1B). Reaction time differences were therefore small but significant at the group level. Three of the 13 subjects showed within subject conflict induced slowing ($p<0.05$, permutation testing), and none of the

subjects showed conflict induced speeding. High conflict trials were also associated with a significant increase in error rate (14 ± 3 vs. 8 ± 3 %, $p < 0.01$, permutation testing). Due to the low number of errors committed by each subject (4.3 ± 1.2 errors for high conflict and 2.5 ± 1.0 errors for low conflict trials), the remaining analyses relate only to correct trials. It is important to emphasize that all changes were gradual and that at no point in the task were subjects informed that a trial had commenced (i.e. there was no "cue"). Therefore, all data presented below are aligned to the response.

4.3.2 Power changes during low and high conflict correct trials

Beginning prior to the response, low conflict trials demonstrated a decrease in beta (16 to 32 Hz) power and an increase in delta (2 to 4 Hz) power in the STN LFP (Figure 4.2A). High conflict trials also demonstrated a decrease in beta power and an increase in delta power in the STN LFP beginning prior to the response (Figure 4.2B). However, the delta power began increasing earlier than it did in the low conflict trials and was associated with a significant increase in theta (4 to 8 Hz) power that was not present in the low conflict trials. Figure 4.2C denotes time-frequency points where LFP power was significantly different between the two trial types. Over the time window from – 3 to -1 s, the mean theta power change relative to the overall power recorded was 10.9 ± 4.3 % for high conflict trials and -1.5 ± 1.5 % for low conflict trials. In the delta band, these values were 5.0 ± 1.8 % for high conflict trials and -3.8 ± 1.4 % for low conflict trials (ANOVA, within-subject repeated measures, conflict x frequency: conflict $F= 16.96$, $df=1$, $p=0.0001$; frequency $F=2.55$, $df=1$, $p=0.11$; interaction $F=0.48$, $df=1$, $p=0.48$). As further reinforced in Figure 4.2D, the changes in theta and delta power in the high conflict trials began at least 2 s prior to the response, when the number of dots moving coherently in one direction was similar to the number of dots moving coherently in the opposite direction. There was no difference in the localization of conflict-triggered delta and theta power increases with respect to

beta power across STN contacts (see section 4.2.6 “Calculation of power and coherence distributions across STN bipolar contacts” above).

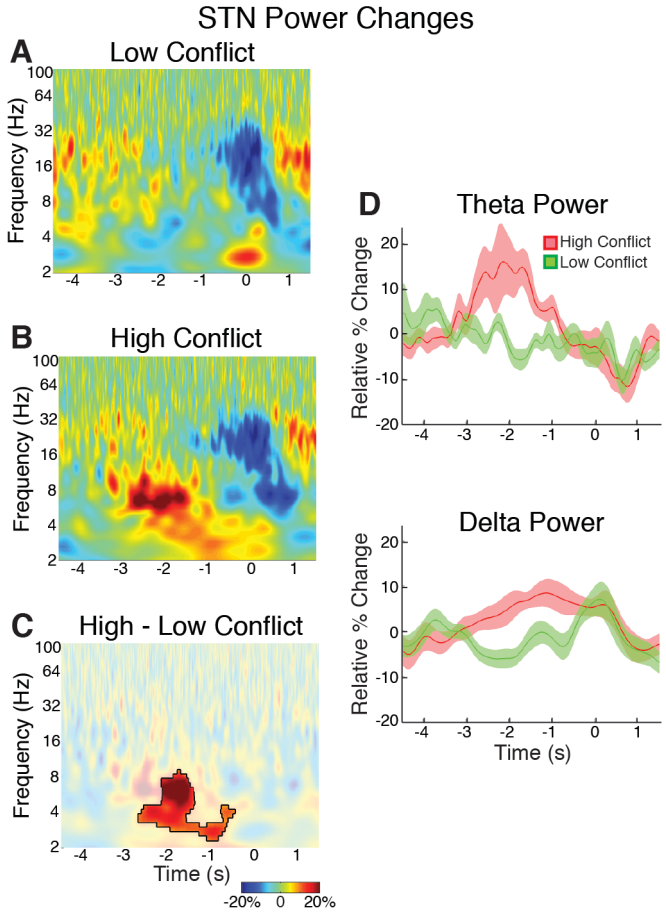


Figure 4.2. Power changes in STN LFP.
(A) Low conflict trials showed a decrease in beta power and an increase in delta power beginning just before response onset. **(B)** High conflict trials showed a response locked theta-delta power increase that began even earlier. The beta power decrease was unchanged. **(C)** Differences were significant (unmasked area, $p<0.05$ after correcting for multiple comparisons, permutation test). **(D)** Theta and Delta power change over time. Mean ± SEM across all STNs (n=23) are shown. T=0 corresponds to response onset. Reproduced with permission from (Zavala et al., 2014).

Bipolar EEG was simultaneously recorded over midline frontal cortex (Fz-Cz) and parietal cortex (Pz-Cz) during the experiment. The only consistent finding was a decrease in beta power that began prior to the response; however, there was no significant conflict-related difference in induced power (data not shown). Conducting the same analysis using the Fz-Pz bipolar configuration also failed to show any significant conflict-related differences in induced power (data not shown).

4.3.3 Coherence differences between EEG and STN in low and high conflict trials

High conflict trials demonstrated significantly higher coherence in the theta-delta band between the STN LFP and frontal EEG relative to the low conflict trials (Figure 4.3A). This coherence difference between trial types was not observed over the parietal cortex (Figure 4.3B). Although response locked, the coherence increase peaked about 2-3 s before the response. As power analyses revealed no difference in the behaviour or localization of subthalamic theta and delta activities over the key time-window of -3 to -2s before the response, and changes in coherence failed to distinguish between these two bands, we will consider the theta-delta band as one for the rest of this chapter. During this key time window, the mean change in theta-delta coherence relative to the average coherence observed throughout the entire experiment was 43.3 ± 20.6 % for high conflict trials and -8.4 ± 8.9 % for low conflict trials ($p<0.01$, permutation testing). The corresponding values for the STN parietal cortex coherence were 21.3 ± 11.8 % for the high conflict trials and 9.5 ± 12.3 % for the low conflict trials, confirming that there was no significant difference in STN parietal coherence increase during conflict ($p>0.05$, permutation testing). Finally, a continuous, time-evolving estimate of the cortico:STN coherence (see methods section 2.2.5 “Inter-site Lachaux coherence”) further confirmed that the conflict related increases in the theta-delta band coherence were specific to the frontal cortex (Figure 4.4). Even when only the consistent nature of the phase difference between the cortex and the STN was considered (See methods section 2.2.4 ”Inter-site phase coherence”), similar results were obtained (Figure 4.4B). There was no difference between localization of theta-delta frequency band coherence with frontal EEG across STN contacts (see section 4.2.6 “Calculation of power and coherence distributions across STN bipolar contacts” above).

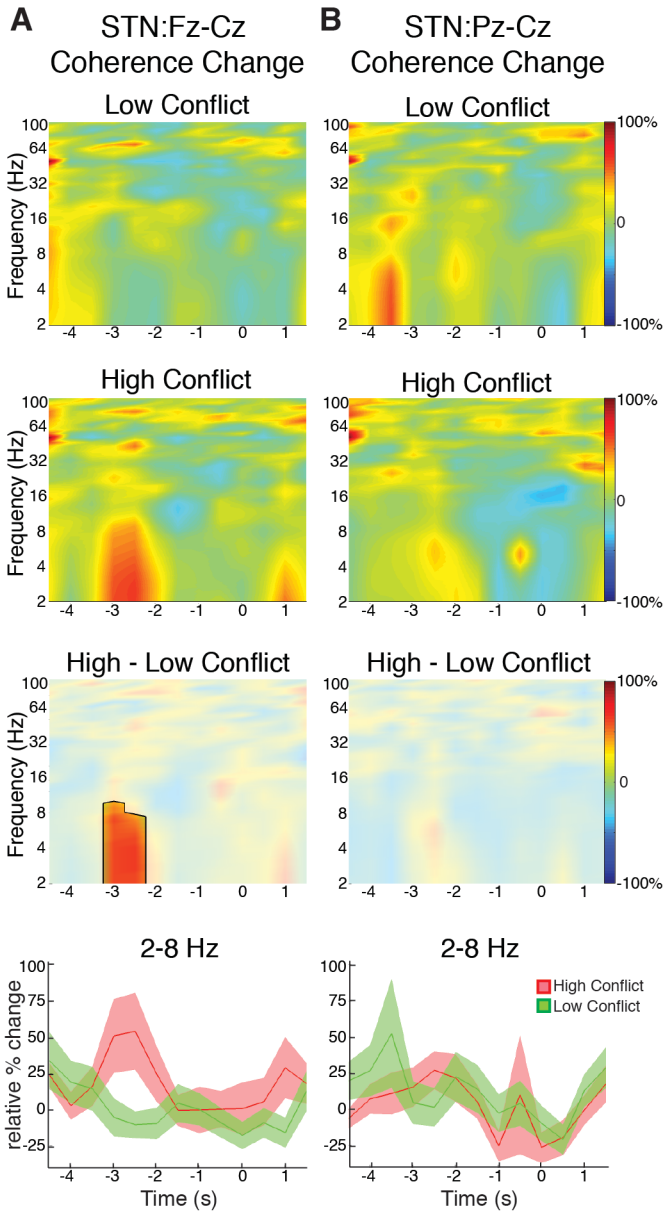


Figure 4.3. Changes in cortico-STN LFP coherence.
(A) High conflict trials showed a relative increase in response locked STN-frontal (Fz-Cz) cortex coherence compared to low conflict trials. Differences were significant (unmasked area, third row, $p<0.05$ after correcting for multiple comparisons, permutation test). **(B)** There were no conflict related changes in STN-parietal (Pz-Cz) cortex coupling. Lowermost panels: Mean ± SEM (n=23) theta-delta coherence changes over time. T = 0 corresponds to response onset. Reproduced with permission from (Zavala et al., 2014).

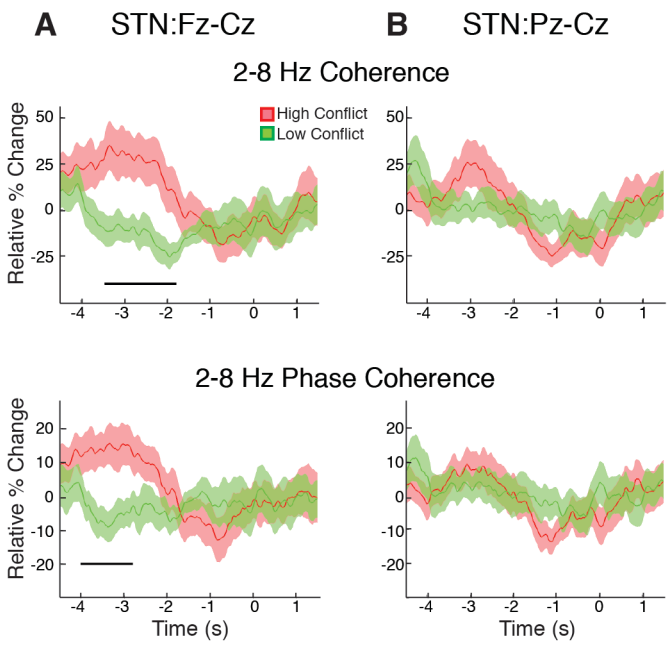


Figure 4.4. Changes in continuous cortico:STN 2-8 Hz coherence using Hilbert transform.
(A, Top) High conflict trials showed a relative increase in response locked STN-frontal (Fz-Cz) cortex coherence compared to low conflict trials. Differences were significant (horizontal bar, $p<0.05$ after comparing for multiple comparisons, permutation test). **(A, Bottom)** Similar results were obtained when only the consistency of the inter-site phase difference was considered. **(B)** There were no conflict related changes in STN-parietal (Pz-Cz) cortex coupling. All panels: Mean ± SEM (n=23) theta-delta coherence changes over time. T = 0 corresponds to response onset. Reproduced with permission from (Zavala et al., 2014).

We then tested whether this increase in low frequency STN-frontal coherence in high conflict trials was symmetrical or biased in one or other direction. Only the frontal cortex-to-STN direction revealed any significant conflict related differences, peaking again at about 2-3 s before the response (Figure 4.5). During this time window, the mean change in frontal cortex to STN Granger causality relative to the average Granger causality observed throughout the entire experiment was 17.9 ± 9.7% for high conflict trials and -10.0 ± 5.2 % for low conflict trials ($p<0.001$, permutation testing). The corresponding values for the STN-to-frontal cortex direction were 5.7 ± 6.3 % for the high conflict trials and -7.7 ± 8.2 % for the low conflict trials confirming that there was no significant STN-to-frontal cortex coherence increase during conflict ($p>0.05$, permutation testing). In order to confirm these results, we repeated the Granger causality analysis using slightly modified parameters (Figure 4.5, model order of 20 with 100 Hz sampling rate, see Methods section 2.2.6 “Granger Causality analysis”). This resulted in a frequency resolution of 5

Hz, and demonstrated qualitatively similar results as those obtained with the BIC determined model order and a 250 Hz sampling rate. In accordance with Figure 4.3B, repeating the Granger causality analysis for the posterior cortex recordings (Pz-Cz) did not reveal any significant differences in cortex-to-STN and STN-to-cortex Granger causality (data not shown).

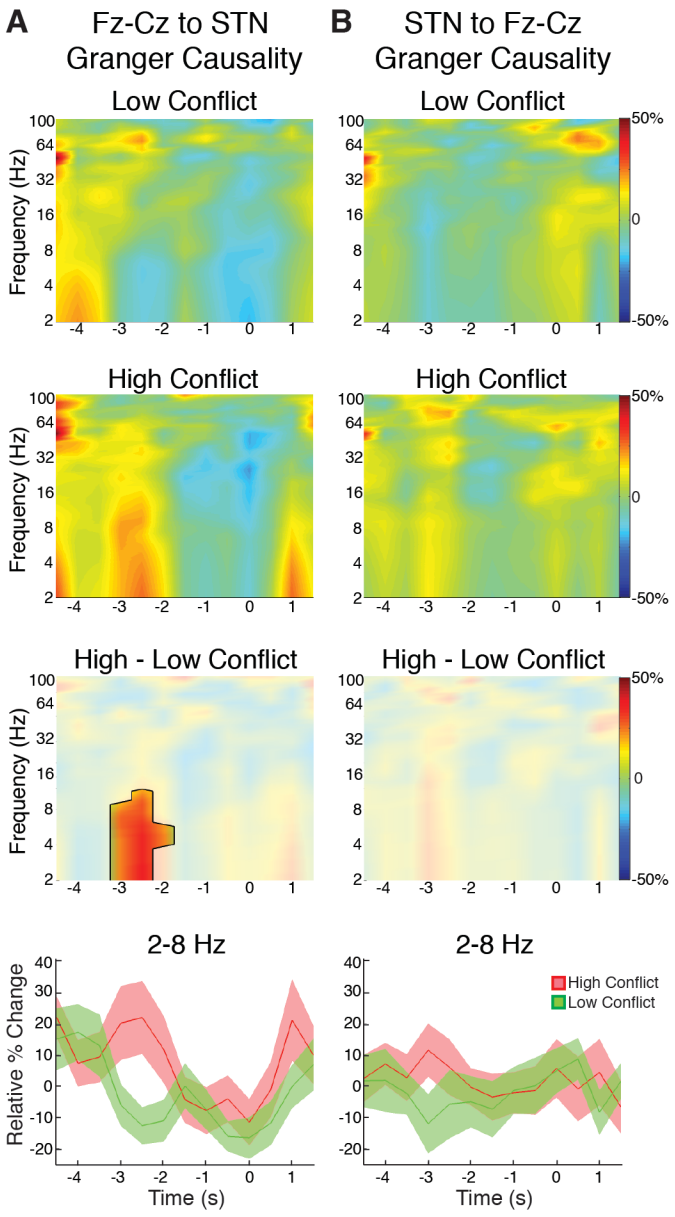


Figure 4.5. Changes in directed coherence between Fz-Cz and STN LFP.
(A) DTF analysis shows a conflict related increase in response aligned directed coherence in the frontal cortex-to-STN direction. Differences between low and high conflict trials were significant (unmasked area, third row, $p<0.05$ after correcting for multiple comparisons, permutation test). **(B)** There were no conflict related changes in response aligned directed coherence in the STN-to-frontal cortex direction. Lowermost panels: Mean ± SEM (n=23) theta-delta directed

coherence changes over time. T = 0 corresponds to response onset. Reproduced with permission from (Zavala et al., 2014).

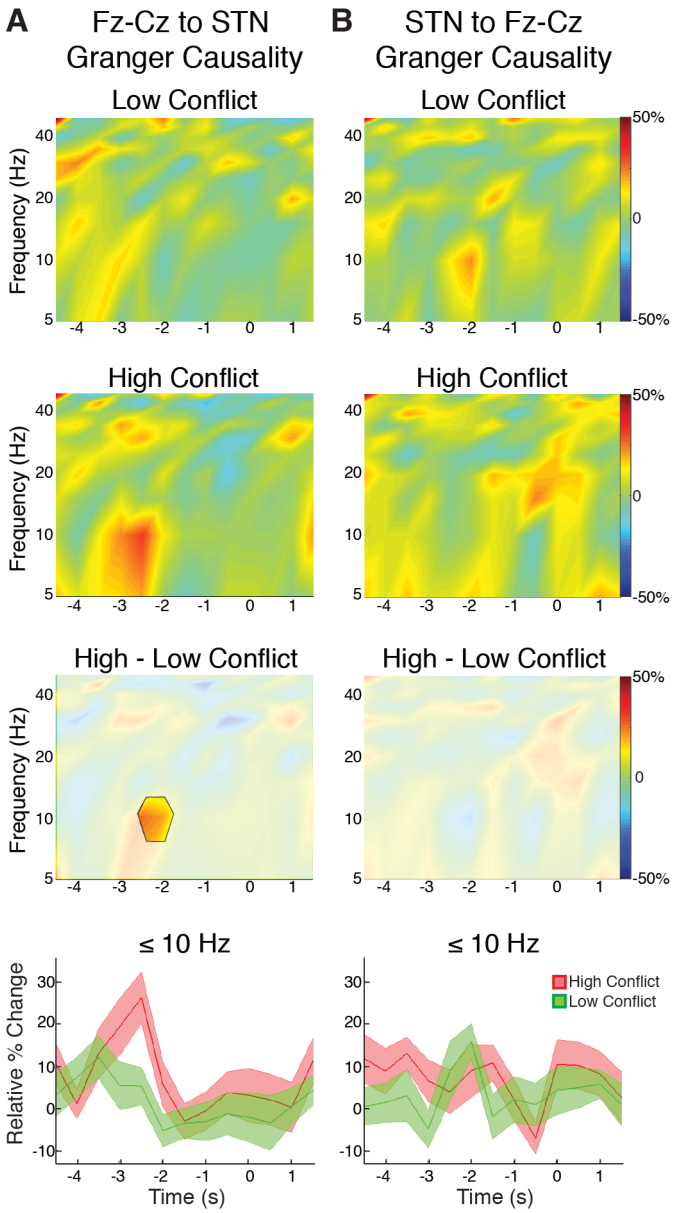


Figure 4.6. Finer frequency resolution directed coherence between Fz-Cz and STN LFP.
Same Granger analysis as that shown in **Figure 4.5** but using a sampling rate of 100 Hz and a model order of 20.

4.4 Discussion

To our knowledge, this is the first study to demonstrate dynamic coupling of neuronal activity between two key brain areas during conflict, the midline frontal cortex and the STN, and to indicate that this coupling is dominated by information flow from cortex to STN. As such, our findings provide critical support for the hypothesis that the STN and the mPFC work together to

raise the brain's decision threshold during times when individuals should "hold their horses" (Frank et al., 2007; Cavanagh et al., 2011; Green et al., 2013). Though prior studies have looked at the behaviour of the STN during conflict (Cavanagh et al., 2011; Fumagalli et al., 2011; Brittain et al., 2012; Zaghloul et al., 2012; Zavala et al., 2013), this study is the first to isolate the STN activity due to conflict from that related either to the processing of stimulus onset or the execution of the response. This was achieved in two ways. First, by using gradual rather than discrete changes in stimuli we avoided a discrete stimulus onset (O'Connell et al., 2012). Second, we arranged for conflict to peak after about 0.8s, which was approximately 2.5 s earlier than the typical response time (about 3.3 s). Theta-delta band activity in the STN and coherence between the STN and the mPFC increased in high conflict trials during this early period of conflict, and it was not until the activity was no longer present that a response was actually executed. Thus, our study confirms that the elevated STN theta activity we and others have observed is indeed specific to conflict and not a signal of time-on-task as others have proposed (Nachev, 2011; Yeung et al., 2011; Cohen and Nigbur, 2013; Scherbaum and Dshemuchadse, 2013; Zavala et al., 2013). It also suggests that the theta power increases observed here (Zavala et al., 2014), in our previous task (Zavala et al., 2013, Chapter 3 above), and in other studies recording from the STN during conflict (Cavanagh et al., 2011; Fumagalli et al., 2011; Brittain et al., 2012) are anti-kinetic in nature.

In contrast, one of the spectral changes associated with response execution (regardless of whether or not conflict was present) was an increase in STN delta power. Recently, Alegre et al. (2013) demonstrated an increase in STN delta power during a stop signal reaction time task. This increase in power was triggered by a go stimulus, but was not significantly different between go trials, stop trials, and failed stop trials. Likewise, our delta power increase at the time of the response was not affected by trial type, suggesting that it is a signal specifically associated with the STN's role in movement execution, and not with the onset of an imperative "Go" stimulus

(Alegre et al., 2013) or of a conflicting flanker or Stroop stimulus (Brittain et al., 2012; Zavala et al., 2013). Thus, despite the traditional rate model of basal ganglia architecture (Albin et al., 1989; DeLong, 1990), increases in STN activity need not be exclusively anti-kinetic (also see section 5.3.3 "Microelectrode single unit firing changes" below).

High conflict trials were also associated with a higher delta power than in low conflict trials, and this began concurrently with the theta increase and extended until about 0.5 s before the response, when the peri-response delta increase seen in low conflict trials set in. One possible explanation for this is that the delta band activity was triggered earlier by the coherently moving dots during the high conflict trials because there were twice as many dots moving coherently than at similar time points in the low conflict trials. According to this interpretation, the elevated theta power observed during conflict may have prevented the delta increases from executing the response. Only when the theta power decreased back to baseline levels was the response executed. It should be kept in mind, however, that we cannot exclude the possibility that some if not all of the early delta power increase represents spectral leakage of the greater power in the theta band during this time period.

4.4.1 Role of theta-delta activity in conflict

The elevated theta-delta power associated with high conflict was accompanied by an increase in frontal cortical driven cortico-STN coherence. The mPFC, which sends projections directly to the STN (Smith et al., 1998; Nambu et al., 2002; Aron et al., 2007; Haynes and Haber, 2013), has long been postulated to be the area of the brain associated with monitoring conflict (Botvinick et al., 2004). Recently, theta band oscillations have received attention as being the potential "language" the mPFC uses to communicate with other brain structures during conflict (Wang et al., 2005; Hanslmayr et al., 2007; Cohen and Cavanagh, 2011; Cavanagh et al., 2012). Crucially, it has been shown that disrupting STN function through high frequency DBS dislocates the relationship between mPFC theta power and reaction time during high conflict scenarios

(Cavanagh et al., 2011). Furthermore, STN DBS weakens the relationship between conflict and mPFC blood oxygen level dependent MRI signalling (Schroeder et al., 2002; Ballanger et al., 2009). In one of the most recent computational models of decision-making, Wieke and Frank (Wiecki and Frank, 2013) propose a simple model in which the anterior cingulate cortex, which is widely believed to be the structure in the mPFC that generates theta oscillations (Gevins et al., 1997; Pizzagalli et al., 2003; Wang et al., 2005) and detects conflict (Botvinick et al., 2004), activates the STN when it perceives that conflicting responses have been triggered by sensory input. Here, we have shown a theta-delta frequency band increase in midline frontal cortical-STN coupling specific to high conflict trials. In the context of the Wiecke-Frank model, this provides compelling evidence of a direct functional link between these two structures and conflict-related processing.

Much of the literature implicating the STN as a crucial player in conflict resolution does so in the context of the drift diffusion model of decision-making (Ratcliff and McKoon, 2008; Cavanagh et al., 2011; Green et al., 2013). The STN, with its ability to inhibit movement via activation of the GPi and with its numerous connections to various sensory and conflict detecting areas of the cortex (Smith et al., 1998; Aron et al., 2007), is hypothesized to adjust the decision threshold during high conflict decisions by inhibiting movement whenever it (or a structure that activates it) detects conflict (Frank, 2006; Bogacz and Gurney, 2007; Cavanagh et al., 2011). Several studies have shown that when STN function is disrupted by DBS, impulsive mistakes are made during conflict presumably because the STN can no longer adjust the threshold according to the amount of conflict present in a trial (Frank et al., 2007; Wylie et al., 2010; Cavanagh et al., 2011; Green et al., 2013). From the point of view of the drift diffusion model, it is tempting to relate the increased theta-delta band drive from the mPFC and correspondingly increased activity in STN seen during conflict with the evidence threshold adjuster of the model.

Taken as a whole, this study yields insight into the major role played by low frequency STN oscillations during conflict processing. Most of the computational and animal models that focus on the STN's role in decision-making, however, are based off of the rate model (see section 1.1.2 "The Basal Ganglia Rate Model: The direct and indirect pathways" and section 1.2 "Computational models of the STN's role in inhibiting choices and actions") and thus consider only single-unit firing (Frank, 2006; Bogacz and Gurney, 2007; Wiecki and Frank, 2013). Therefore, one important next step was to test how the oscillatory changes we report here during conflict interact with the firing rate changes that have been observed in the STN during conflict and response inhibition (Isoda and Hikosaka, 2008; Zaghloul et al., 2012; Schmidt et al., 2013).

Chapter 5. Human subthalamic nucleus theta and beta oscillations entrain neuronal firing during sensorimotor conflict

5.1 Introduction

Our previous work has demonstrated the importance of STN oscillations, particularly those in the theta band. Increases in STN theta power (Cavanagh et al., 2011; Fumagalli et al., 2011; Brittain et al., 2012; Zavala et al., 2013, 2014) are similar to those observed over the medial prefrontal cortex (mPFC) (Cohen and Cavanagh, 2011; Cavanagh et al., 2012). These results, together with the results presented in Chapter 4 (Zavala et al., 2014), suggest that coherent theta oscillations may be the neural substrate by which conflict related information is conveyed between these structures. Traditionally, however, it is the firing rate dynamics of the basal ganglia that are assumed to be responsible for determining whether or not a response is executed (Albin et al., 1989; DeLong, 1990). Accordingly, several studies have demonstrated that STN spiking activity increases during response inhibition paradigms to prevent or delay responses (Isoda and Hikosaka, 2008; Zaghloul et al., 2012; Burbaud et al., 2013; Schmidt et al., 2013; Bastin et al., 2014). According to one study (Schmidt et al., 2013), response inhibition involves a race between pro-kinetic inhibition of the GPi via cortico-striatal inputs and the anti-kinetic excitation of the GPi via STN firing (see Figure 1.1).

As both STN oscillatory power and spiking activity are related to decision-making, an important yet unanswered question is how to reconcile the classic neuronal firing rate model (Albin et al., 1989; DeLong, 1990) of basal ganglia architecture with the role played by STN neuronal oscillations as suggested by our previous work. This question has remained elusive partially because most studies involving behaviorally relevant recordings from the human STN are usually restricted to either LFP or single unit recordings alone. Here we directly studied these neural

mechanisms by examining simultaneously recorded LFP and spiking activity from the human STN as participants undergoing DBS surgery for Parkinson's disease performed an Eriksen flanker task. We identify separate populations of spiking neurons in the STN, and specifically examined interactions between their spiking activity and changes in theta and beta oscillations observed in the LFP.

5.2 Methods

5.2.1 Intraoperative recordings during deep brain stimulation surgery

Intraoperative recordings were captured from the STN in participants undergoing DBS surgery of the STN for Parkinson's disease. Patient details are available in Table 5.1. The study was conducted in accordance with a NIH IRB approved protocol, and all subjects gave their written informed consent to take part in the study. Subjects received no financial compensation for their participation. Parkinson's medications were stopped on the night before surgery (12 h preoperatively). Recordings were captured while participants were alert, at rest, and in an OFF state while in the operating room.

Table 5.1. Clinical details from second Flanker experiment.

Case	Age	Disease Duration	UPDRS Off (III)	UPDRS On (III)	Reasons for Surgery
1	69	13	33	18	Bradykinesia
2	46	12	31	20	Rigidity & bradykinesia
3	71	6	24	12	Tremor & bradykinesia
4	69	14	38	22	Rigidity & bradykinesia
5	53	35	46	32	Bradykinesia & rigidity
6	64	10	34	28	Tremor & rigidity & bradykinesia
7	64	20	49	21	Bradykinesia
8	49	7	25	12	Bradykinesia
9	42	5	50	14	Tremor & bradykinesia
10	63	20	29	16	Tremor & bradykinesia

11	53	13	38	22	Tremor & bradykinesia
12	48	9	48	12	Rigidity & bradykinesia
13	69	14	51	24	Bradykinesia & rigidity
14	47	14	40	11	Tremor & bradykinesia
15	65	15	41	20	Tremor & bradykinesia

5.2.2 Flanker task

Participants performed an Eriksen flanker sensorimotor task (Figure 5.1A) (Eriksen and Eriksen, 1974) in the intra-operative environment on a testing laptop (PyEPL, Geller et al., 2007). On each trial, a white fixation dot initially appeared in the centre of the screen for 900 ms. Flanking arrows, all pointing in the same leftward or rightward direction, were then subsequently displayed to the left and right of the fixation dot. Following the flanking arrows by 100 ms, a target arrow was presented between the flanking arrows above the fixation dot. Target arrows could point in the same direction (low conflict) or in the opposite direction (high conflict) as the flanking arrows. Arrows were displayed on the screen until the participant indicated the direction of the target arrow using a leftward or rightward movement of a joystick, or until four seconds elapsed. Participants made all movements with the hand contralateral to the side of intra-operative recording for that session. Trials were separated by a blank screen displayed for a random duration of 1000 ±100 ms. Throughout an experimental session, the number of low and high conflict trials were balanced as were the number of trials requiring leftward and rightward movements.

Participants completed three blocks of 40 trials each during each experimental session. On the day before surgery, participants also performed a complete session in order to familiarize themselves with the task. Three months after surgery, 14 participants performed two additional sessions of the task, first with their DBS stimulators turned on and then with their stimulators

turned off. These post-operative sessions were performed while patients were on their regular Parkinson's medications.

All incorrect trials as well as all trials with response times less than 300 ms and greater than 1500 ms were discarded. In order to exclude participants who were not performing the task correctly, the inverse of the response times was used to normalize their distributions (Carpenter and Williams, 1995), and then an unpaired t-test was used to compare the inverse response times between low and high conflict trials. All experimental sessions that did not exhibit a significant difference between trial types were excluded. Most participants performed one session each while recordings were captured from the left and right STN. Six participants did not complete the second session because of fatigue. Two participants performed an additional session during recordings from one side. For these two participants, the session exhibiting the faster response times was used for LFP power and phase analysis. Because spike firing rate and phase analyses were performed on individual cells, cells from both sessions were included for spiking analysis.

5.2.3 Electrophysiology

All analyses were performed using custom MATLAB scripts (Mathworks, Natick, MA). LFP activity was extracted by bandpass filtering macroelectrode signals between 1 and 500 Hz, notch filtering at 60 Hz, and down-sampling all LFP signals to 1 kHz. Macroelectrode signals were then referenced to the common average of all simultaneously recorded macroelectrode contacts to yield three referenced monopolar channels per STN. Repeating the LFP macroelectrode power and phase analysis using a bipolar configuration produced the same results as those presented in section 5.3.2 “Macroelectrode LFP power and phase changes” of this chapter.

In order to obtain magnitude and instantaneous phase information in the frequency domain, the LFP signals were convolved with complex valued Morlet wavelets (wave number 6). 47 logarithmically spaced (8 scales/octave) wavelets between 2 and 107 Hz were used. For cue-

locked analyses, LFP signals were analysed from 500 ms before to 1500 ms following target arrow presentation. For response-locked analyses, LFP signals were analysed from 1000 ms before to 1000 ms after the response. A 1500 ms buffer was used on both sides of the clipped data to eliminate edge effects. The magnitude of the continuous-time wavelet transform was then squared and log transformed to generate a continuous measure of instantaneous power. The power from each channel was subsequently z-scored separately for each frequency using the mean and standard deviation of the power recorded from that channel during all fixation periods (baseline; 1000 to 200 ms before target arrow onset). The z-scored power from the three channels was then averaged together to avoid selection bias, although averaging across LFP channels within each STN may underestimate spectral changes as not all macroelectrodes were within or bridging the STN. Prior to any subsequent analysis, all trials that exhibited a clear artefact in the LFP trace were discarded.

Spiking activity was extracted by bandpass filtering microelectrode recordings between 0.3 and 3 kHz and resampling the filtered signals at 24 kHz. Individual spike clusters were identified for each recording using principal component analysis (Plexon Offline Sorter, Inc., Dallas, TX). A spike cluster was classified as belonging to a single neuron if < 2% of the spike events in that cluster fell within a refractory period of 1.5 ms and if all spike waveforms clustered together in PCA space. Out of the 32 microelectrode recordings that demonstrated spiking activity, six recordings contained two spike clusters that separated in the PCA space. Data from each of these six recordings were exported as two separate clusters. Henceforth, we will refer to the 38 exported clusters as cells.

5.2.4 LFP power and phase analysis

To assess differences in spectral power between high and low conflict trials across participants, the mean difference in z-scored power were first calculated for each time point and for each frequency separately across all low and high conflict trials for each experimental session and then

averaged across STN recording sessions. A p-value was assigned to each time-frequency point by comparing the true difference to the distribution of differences from 1000 surrogate data sets generated by permuting the low and high conflict trials prior to averaging across trials and sessions. The p-values were then corrected for multiple comparisons using exceedence mass testing (Maris and Oostenveld, 2007), see methods section 2.3.2 “Permutations and exceedance mass testing”.

To assess for differences in ITPC, the differences in ITPC between low and high conflict trials were calculated at every time point and frequency for each channel, and then averaged across channels within each STN before averaging across STNs. The p-value at each point was calculated using the distribution of 1000 permuted values and corrected for multiple comparisons at a significance level of 0.05 using exceedence mass testing (Maris and Oostenveld, 2007), see methods section 2.3.2 “Permutations and exceedance mass testing”.

To determine whether the theta band time series exhibited significant differences between low and high conflict trials, LFP signals were band pass filtered between 5 and 7 Hz, Hilbert transformed, then processed using the same procedure that was applied to the continuous-time wavelet data described above. The results were corrected for multiple comparisons using the same exceedence mass testing procedure described above, where clusters were determined only by time points in the instantaneous time series.

5.2.5 Spiking activity analysis

To determine whether an individual cell was responsive during the task, its average firing rate was calculated during three non-overlapping temporal epochs. We divided the period between target arrow onset and response into three equally sized temporal epochs for each trial based on that trial's response time. We additionally defined a baseline epoch during the fixation period (1000 ms to 200 ms before target arrow onset). Spiking activity during each of the three temporal

epochs of a given trial was then compared to the activity during the baseline epoch of that trial, and significance was assed across all trials using a paired t-test. If any of the three temporal epochs exhibited a significant difference in activity ($p < 0.016$, Bonferroni corrected across epochs), the cell was considered to be responsive. Responsiveness was assessed using three separate temporal epochs because some cells only exhibited brief changes in firing that would otherwise be missed if averaged over the entire decision period. Because previous studies have demonstrated that cells within the STN may preferentially adjust their activity with different movements (Williams et al., 2005; Hanson et al., 2012), we investigated whether this activity was present in our data. To determine if a cell was directionally tuned, trials were divided into rightward and leftward movements of a joystick, and spiking activity was separately compared in each temporal epoch between both movements. Any cell that exhibited a significant direction-related difference in spiking activity in any of the three temporal epochs ($p < 0.016$, Bonferroni corrected) was considered to be directionally tuned.

In order to identify any consistencies in the temporal firing rate dynamics observed during the task across cells, the continuous-time firing rates for each responsive cell were calculated by smoothing the spike train from each trial (1 ms bins) with a Gaussian kernel (standard deviation 100 ms). To generate a z-scored firing rate, the continuous time firing rates for each trial were compared to the mean and standard deviation of continuous time firing rates during the baseline period (1000 to 200 ms before target arrow onset) and then averaged across trials. To determine if different cells clustered into populations that exhibited similar temporal dynamics, principal component analysis (PCA) was used on the resulting mean firing rates (Narayanan and Laubach, 2009), and based on the first two components, the cells were divided into two separate populations using k-means clustering.

To determine if continuous-time z-scored firing rates exhibited a significant difference between high and low conflict trials across cells, the mean difference in firing rate between trial types was

calculated for every cell at every time point. The average difference was then calculated across all cells within a given group (early cells, late cells, down cells) and the difference was compared to an empiric distribution generated by permuting the low and high conflict trial labels 1000 times. It was also determined whether differences in spike histograms were significant between trial types by first binning spiking data into overlapping 100 ms bins (25 ms overlap), z-scoring the firing rate by the baseline period bins, and calculating the mean difference in spiking activity between low and high conflict trials for each bin. To assess differences in the rate of change of spiking activity, the first derivative of the continuous firing rates between trial types was compared. Conflict related differences in all three of these metrics were corrected for multiple comparisons using exceedence mass testing (Maris and Oostenveld, 2007), see methods section 2.3.2 "Permutations and exceedance mass testing". Time points during which the firing rate was significantly ($p < 0.05$) greater than zero across cells were also identified by using a single sample t-test. The difference in onset (or offset) time between low and high conflict trials was also calculated and compared to the distribution of 1000 onset (or offset) differences generated by permuting the low and high conflict trials prior to finding the onset (or offset) time.

5.2.6 Spike-phase interactions

Spike phase interactions were calculated for each cell by first filtering LFP signals within each frequency band (delta 2-4 Hz; theta 5-7 Hz; alpha 8-14 Hz; beta 15-30 Hz) and then extracting instantaneous phase information using the Hilbert transform. The instantaneous phase of the macroelectrode LFP signal was stored for all spiking events captured on the associated microelectrode between the stimulus onset and the response. The mean vector length of phases was then separately calculated across all spike events in all low and high conflict trials similar to the way ITPC is calculated. This yielded a separate spike phase locking value for each trial type (r_{low} and r_{high}). The resulting spike phase locking values were then normalized by permuting phase

information across trials. Whereas in the true case, spike times from a given trial were assigned an instantaneous phase from that same time point in the same trial, in each permuted case each spike time was assigned an instantaneous phase from the same time point in a different trial drawn at random from the pool of similar conflict level trials (i.e. r_{low} values were normalized by low conflict trials and r_{high} values were normalized by high conflict trials). The phase information was permuted separately for low and high conflict trials to prevent any biases that may emerge from non-uniform phase distributions caused by differences in ITPC. In this manner, spike phase locking values observed for a given trial type were normalized by the probability of observing spike phase locking by chance given the distribution of phases and spike times observed during that trial type. The phase information was permuted 1000 times resulting in a distribution of 1000 surrogate r_{low} and r_{high} values for each cell. The true r_{low} and r_{high} were then compared to the mean and standard deviation of the distribution of permuted values to generate a normalized spike phase locking value for each cell and for each trial type (R_{low} and R_{high}). To test for differences in spike phase locking between trial types, the distribution of normalized values was compared between low and high conflict trials across cells using a paired t-test. To generated a continuous-metric of spike phase-locking, the normalized mean vector length for each trial type was calculated during sliding 750 ms windows, stepped every 50 ms from 500 ms before to 1500 ms following cue presentation, and from 1000 ms before to 1000 ms following response. To identify contiguous time points that showed significant conflict related differences while correcting for multiple comparisons, the permutation and exceedence mass testing procedures described above were used, where clusters were determined only by time points in the spike phase locking time series.

To confirm the observed differences in spike phase locking between low and high conflict trials, the analysis was repeated using the pairwise phase consistency (PPC, Vinck et al., 2010) in the delta, theta, alpha, and beta band. The instantaneous phase of each spiking event was extracted

and the cosine similarity between phases associated with all pairs of spiking events was calculated for the low (PPC_{low}) and high (PPC_{high}) conflict trials separately. The resulting PPC_{low} and PPC_{high} values were then normalized separately by using an identical permutation procedure as the one used to normalize the spike phase locking r values. To further confirm the conflict related differences we report, the spike phase locking was recalculated using the instantaneous phase of the continuous time wavelet transform for each frequency from 2 to 107 Hz during each trial. The mean vector length of phases across all spike events was calculated for each frequency, and then normalized separately for low and high conflict trials as described above to generate a normalized spike phase locking value for each frequency, cell, and trial type. For every frequency, the normalized spike phase locking values were averaged across cells to calculated differences between low and high conflict trials. To identify contiguous frequencies that showed significant conflict related differences while correcting for multiple comparisons, the permutation and exceedence mass testing procedures described above were used, where clusters were determined only by frequencies in the spike phase locking spectrum.

To visualize spike-phase interactions in a single cell (Figure 5.8), the Hilbert transform was used to extract the instantaneous phase information of the filtered theta (5-7 Hz) or beta (15-30 Hz) LFP signal captured from the corresponding macroelectrode. All of the troughs of the oscillation observed between the target arrow onset and the motor response were identified in each high conflict trial. The continuous-time firing rate was then calculated in a 600 ms (theta) or 400 ms (beta) temporal epoch surrounding each trough by convolving the spike train (1 ms bins) with a Gaussian kernel (standard deviation 15 ms). A 15 ms kernel was used in order to maintain the temporal resolution necessary to discern spike-phase interactions. The continuous time firing rates during each epoch were detrended using matlab's 'detrend' function and then averaged the across all epochs. To visualize theta and beta oscillations, the raw macroelectrode LFP trace was also extracted in the same epochs and averaged across all epochs during all high conflict trials.

5.3 Results

5.3.1 Task and Behavior

Intraoperative micro and macroelectrode recordings from the STN were captured during DBS surgery for Parkinson's disease in 15 participants (5 male; 58.1 ± 2.67 (mean ± SEM) years old) as they performed an Eriksen flanker task (Figure 5.1A). Because DBS electrodes were implanted bilaterally in all participants, some participants performed the task twice yielding 26 total - experimental sessions. During each session, recordings were made from three macro/microelectrode pairs to investigate local field potential (LFP) and single-unit spiking activity in the STN (Figure 5.1A).

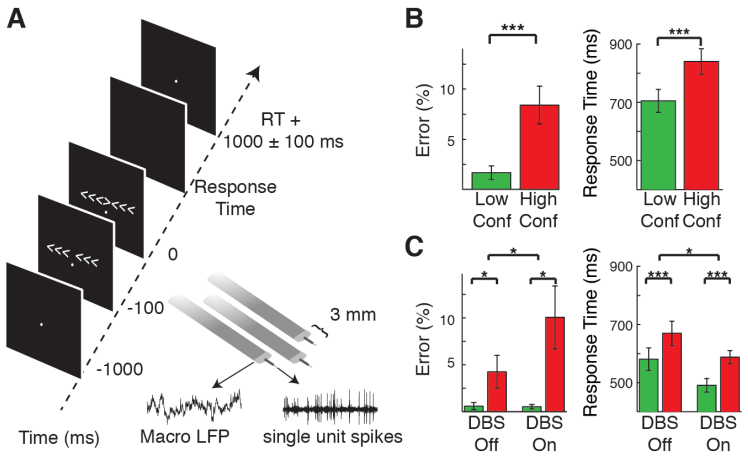


Figure 5.1. Flanker task during intraoperative STN recordings.
(A) During each trial, a warning cue appeared 900 ms before the onset of flanking arrows. A central arrow appeared 100 ms later, and all arrows remained on the screen until a response was made. Central arrows could be congruent (low conflict) or incongruent (high conflict) with flanking arrows. Inset, three macro/microelectrode pairs captured LFP (macro) and single unit spiking activity (micro) from the STN as participants performed the task. **(B)** Error rates during the intraoperative sessions for high and low conflict trials, averaged across participants (mean ± SEM; *** denotes $p < 0.001$), are shown on the left. Average response times for the correct high and low conflict trials are shown on the right. **(C)** Error rates during postoperative sessions for high and low conflict trials with DBS On and Off (mean ± SEM; *** denotes $p < 0.001$; * denotes $p<0.05$) are shown on the left. Response times during high conflict trials with DBS On and Off are shown on the right. Only the error rates showed a significant conflict x stimulation interaction.

In 19 of the experimental sessions (13 participants), participants exhibited significantly greater response times ($p < 0.05$, unpaired t-test) during high conflict trials compared to low conflict trials. In these sessions, participants were significantly more likely to make a correct movement during low conflict (98 ± 1%) compared to high conflict trials (91 ± 2%; $t(18) = -4.36$, $p = 0.001$;

Figure 5.1B). We retained correct trials in these sessions for subsequent analyses. During the retained sessions, participants responded with a joystick movement 840.10 ± 43.83 ms following cue presentation during high conflict trials and 704.98 ± 39.65 ms following cue presentation during low conflict trials ($t(18) = 11.81$, $p = 0.001$; Figure 5.1B).

Fourteen of the participants performed the task two additional times, once with the implanted DBS stimulator turned on and once with the stimulator turned off, during an outpatient follow up three months following surgery. These participants exhibited significantly higher error rates during high conflict trials, significantly more errors when their stimulators were on, and a significant interaction between stimulation and conflict (ANOVA, within-subject repeated measures, conflict x stimulation: conflict $F(1,12)= 7.74$, $p=0.016$; stimulation $F(1,12)=7.37$, $p=0.018$; interaction $F(1,12)=9.08$, $p=0.01$, Figure 5.1C). Reaction time differences also showed a main effect of conflict and stimulation (conflict $F(1,12)= 101.94$, $p<0.001$; stimulation $F(1,12)=14.67$, $p=0.002$), but there was not a significant interaction between the two (interaction $F(1,12)=1.2$, $p=0.29$).

5.3.2 Macroelectrode LFP power and phase changes

We investigated changes in spectral power in the macroelectrode LFP signal, and found that both low and high conflict trials exhibited similar decreases in beta and increases in gamma band oscillatory power during the task (Figure 5.2A and Figure 5.3A). Post response increases in beta were delayed during high conflict, consistent with the longer response times during these trials. In contrast, the theta band showed significantly greater power during the high conflict trials compared to low conflict trials early in the decision period, prior to the response ($p < 0.05$, permutation test). We confirmed that these changes reflected induced theta activity, rather than changes in evoked event related potentials (see methods section 2.2.2 “Oscillatory power analysis”). We also calculated the macroelectrode LFP ITPC for every frequency and found that

high conflict trials had significantly greater theta band phase consistency relative to low conflict trials following cue presentation ($p < 0.05$, permutation test; Figure 5.2B).

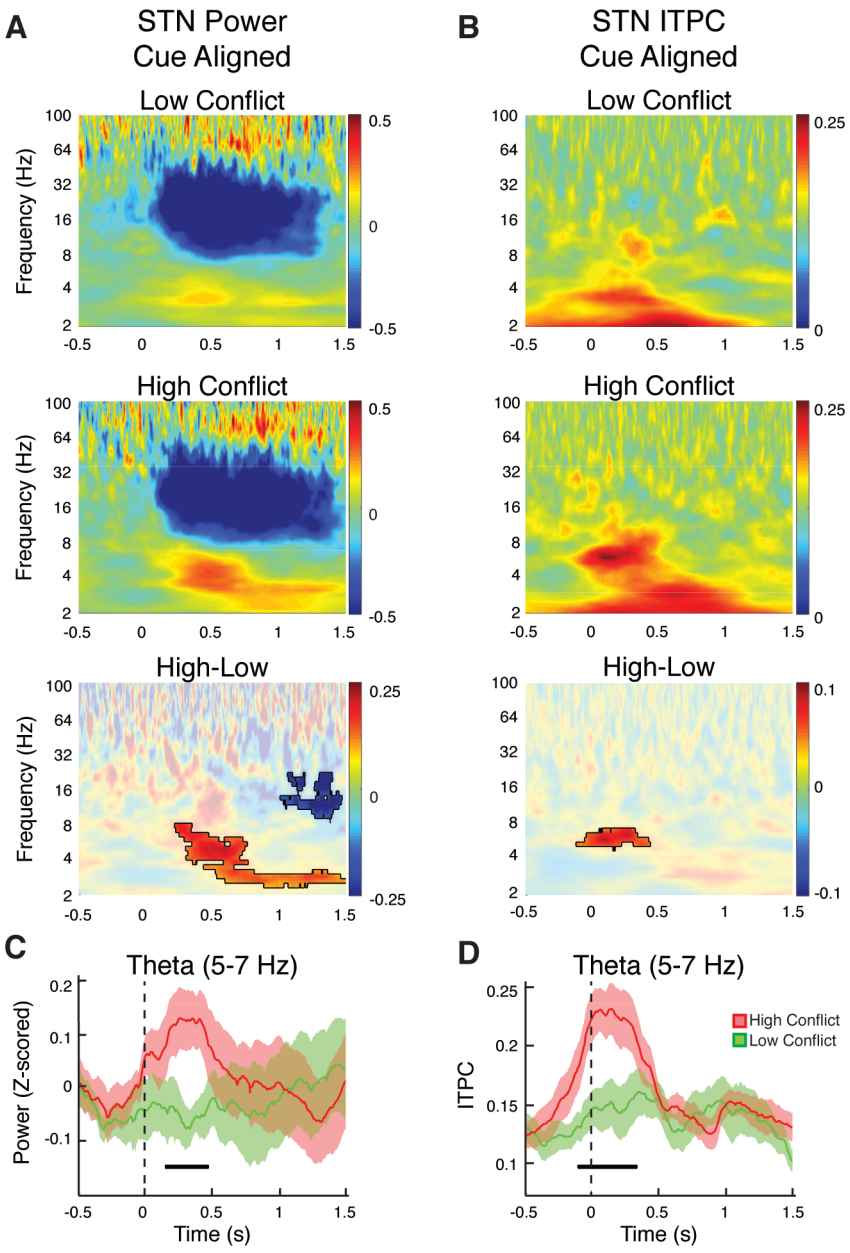


Figure 5.2. STN theta band activity is higher during conflict.
(A) Cue locked z-scored oscillatory power during all correct trials averaged across all STNs (n=19, target arrow presented at t=0). Low conflict trials are shown on the top, high conflict trials are shown in the middle, and the average difference in oscillatory power between high and low conflict trials (masked at $p < 0.05$ after correcting for multiple comparisons, permutation test) is shown on the bottom. Average theta power was significantly higher during high conflict trials. **(B)** Cue locked ITPC during all correct trials averaged across all STNs (n=19, target arrow presented at t=0). Order of plots same as in (A). Average theta (5-7 Hz) ITPC was significantly higher during high conflict trials. **(C)** Average (mean ± SEM) theta (5-7 Hz) power for low (green) and high (red) conflict trials. Time points exhibiting a significant difference ($p<0.05$ after correcting for multiple comparisons, permutation test) between low and high

conflict trials are denoted by the black horizontal bar. **(D)** Average (mean ± SEM) theta (5-7 Hz) ITPC for low (green) and high (red) conflict trials. Significance denoted as in (C).

To explore the time course of theta band activity during low and high conflict trials, we examined the instantaneous power and phase consistency of the LFP signal, bandpass filtered in the frequency range (5-7 Hz) that showed significant differences in power and ITPC between trial types. The cue-locked, z-scored theta power increase was significantly higher in the high conflict trials compared to low conflict trials from 160 to 460 ms following cue presentation ($p < 0.05$, permutation test; Figure 5.2C; for data time-locked to response, see Figure 5.3A,C). In addition, the theta band ITPC was significantly higher in the high conflict trials compared to low conflict trials from -92 to 325 ms relative to the stimulus onset ($p < 0.05$, permutation test; Figure 5.2D). We did not find significant differences in theta phase consistency between high and low conflict trials when we time-locked trials to the motor response (Figure 5.3B,D), indicating that the observed changes in phase consistency were specific to the cue as argued previously in Chapter 3 (Zavala et al., 2013). In disagreement with Chapter 3, however, is the finding that the high conflict trials showed higher theta ITPC. Though our previous work had shown no difference in theta band ITPC between low and high conflict trials, we did previously show that the slowest high conflict trials showed decreased levels of cue-locked theta ITPC (see Chapter 3 above). We interpreted those changes as being due to conflict, however, the findings from the study presented here seem to disagree with this claim. Indeed, looking at the theta band ITPC separately for the fastest and slowest half of the high conflict trials did not reveal any significant differences between the two trial types (data not shown).

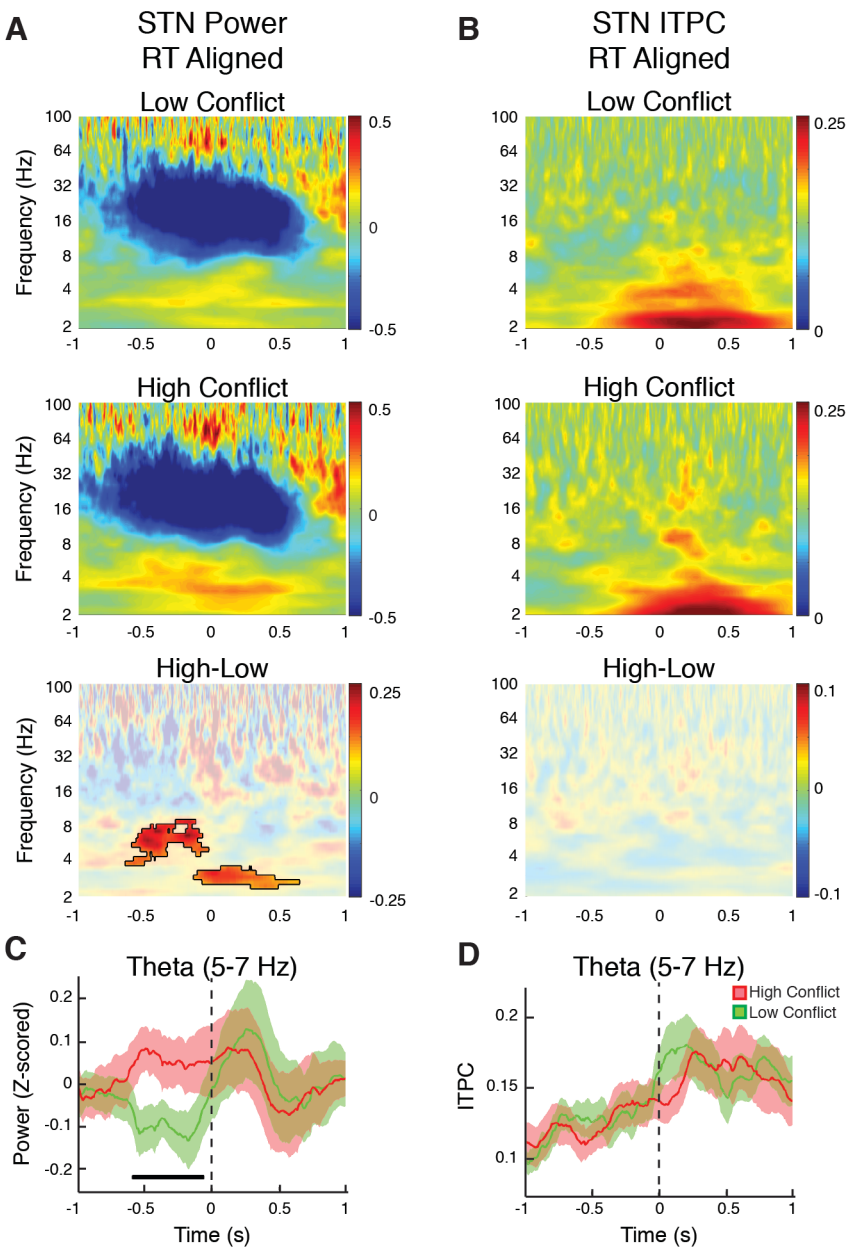


Figure 5.3. Response-aligned STN LFP activity.
Same data from **Figure 5.2**, but aligned such that t=0 corresponds to the response. Significant conflict related differences in oscillatory power in the theta band occurred prior to the response ($p < 0.05$ after correcting for multiple comparisons, permutation test). We also observed a post response increase in delta band power during conflict. The response aligned ITPC data did not demonstrate any significant differences between trial types, indicating that the observed phase consistency changes reported in **Figure 5.2** were specific to the cue.

5.3.3 Microelectrode single unit firing changes

We extracted spiking activity from microelectrode recordings in the STN while participants engaged in the decision task and identified 38 single units, which we will henceforth refer to as cells (33.7 ± 3.2 spikes/s). We examined the spiking activity following cue presentation and

found 22 cells that individually exhibited a significant ($p < 0.016$, paired t-test) increase in spiking activity between cue presentation and the response. Nine of these cells exhibited significant differences in firing rates that depended on the direction of movement of the contralateral hand. Seven additional cells responded with a significant ($p < 0.016$) decrease in firing during the decision period, and nine cells did not demonstrate a significant change. We focused our subsequent analyses on the 29 cells exhibiting a significant change in spiking activity during the task.

Representative spiking activity from a single cell exhibiting a significant increase in firing during the decision period is shown in Figure 5.4A. Spiking activity for this cell is time-locked to the stimulus onset. Representative spiking activity from a different cell is shown in Figure 5.4B and exhibits a peak in firing during the motor response. In order to examine whether these two cells represent two separate populations of cells that exhibit increased activity during the decision period, we used principal component analysis to investigate the temporal dynamics of the spiking response (Narayanan and Laubach, 2009). We identified two principal components of the continuous time firing rate that accounted for at least 10% of the variance across all 22 of the cells that exhibited a significant increase in spiking activity during the task (Figure 5.4C). One principal component accounted for 45% of the variance across cells and exhibited an increase that peaked around the response, while a second component accounted for 34% of the variance and exhibited a peak that occurred immediately following cue presentation. We used k-means clustering to divide these 22 cells into two separate populations based on these first two principal components (Figure 5.4C). This procedure identified one group of cells ($n = 9$) that, on average, exhibited spiking activity that peaked immediately following cue presentation (early cells), and a second group of cells ($n = 13$) that, on average, exhibited spiking activity that peaked around the motor response (late cells; Figure 5.4D,E). Notably, the first principal component alone was sufficient to separate these two populations of cells (Figure 5.4C).

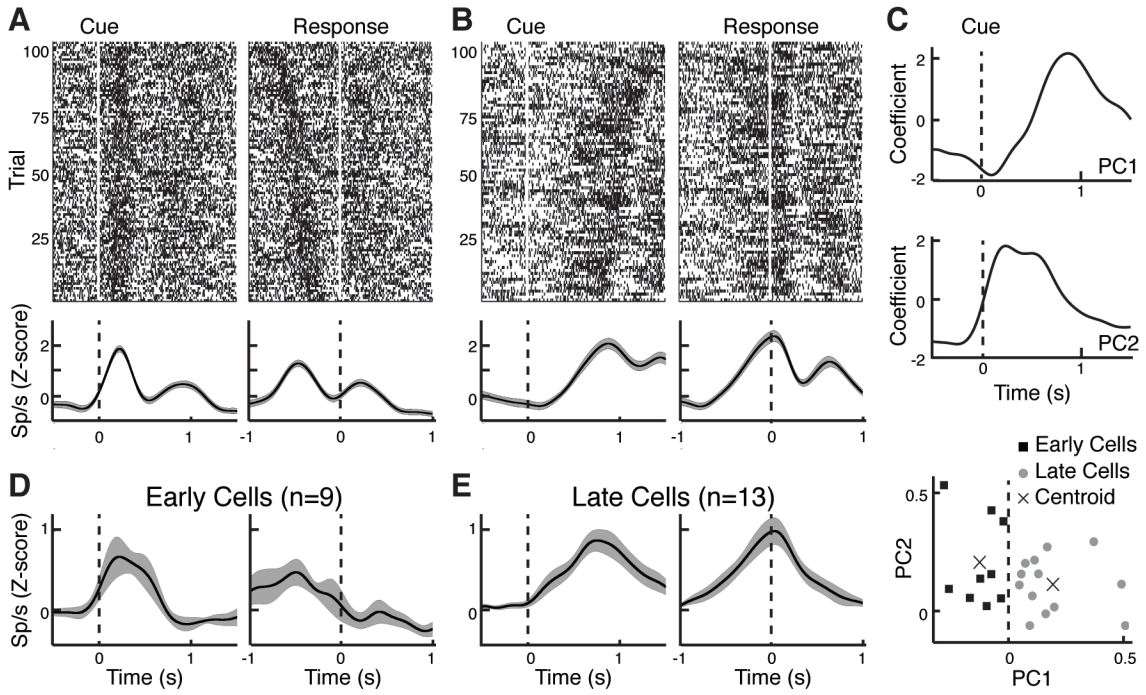


Figure 5.4. Task related increased spiking activity in STN cells.
(A) A representative cell exhibiting increased spiking activity following presentation of the target arrow, locked to cue (left, t=0) and to the response (right, t=0). Raster plot (top) and the average continuous time firing rate (bottom, mean ± SEM) are shown for the cell. Trials are sorted by reaction time. **(B)** A representative cell exhibiting increased spiking activity with the response. **(C)** First two principal components of the firing rate response curves that accounted for over 10% of the variance for all cells that exhibited a significant increase in spiking activity during the decision period (n = 22). The eigenvector for the first component accounted for 45% of the variance (top) and the eigenvector for the second component accounted for 34% of the variance (middle). All cells exhibiting a significant increase in spiking activity clustered into two separate populations based on the loadings of the first two principal components (bottom; black squares represent early cells (n = 9); grey circles represent late cells (n = 13); centroids of each cluster are indicated by an X). **(D)** Average continuous time firing rate (mean ± SEM) for all early cells (n = 9) locked to cue (left) and to the response (right). **(E)** Average continuous time firing rate (mean ± SEM) for all late cells (n = 13) locked to cue (left) and to the response (right).

To explore the effect of conflict on spiking activity, we separately examined spiking activity during high and low conflict trials for early and late cells. Spiking activity for a representative early cell demonstrated an initial increase during both high and low conflict trials that returned to baseline prior to the response. This increase was maintained for a longer period of time during high conflict trials (Figure 5.5A). We found that across the population of early cells, this distinction between high and low conflict trials was consistent and significant ($p < 0.05$, permutation test; Figure 5.5B, top). We also collected spiking activity in 100 ms bins (25 ms overlap) for each trial, and observed that, relative to high conflict trials, low conflict trials had significantly less spiking activity in the contiguous bins from 500 to 650 ms following stimulus

onset ($p < 0.05$, permutation test; Figure 5.5B, bottom). When we calculated the time points at which spiking activity was significantly greater than baseline across cells ($p<0.05$, one sample t-test), we found no significant differences in spiking onset times ($p>0.05$, permutation test). There was a significant difference, however in spiking offset time; the low conflict trials exhibited spiking offset 471 ms following cue presentation and the high conflict trials exhibited spiking activity that returned to baseline 680 ms after the cue ($p<0.05$, permutation test). We observed a similar difference in spiking offset between trial types when locked to the response (low conflict 441 ms and high conflict 177 ms before the response; $p < 0.05$, permutation test). When we analyzed the continuous time firing rates for the early cells aligned to the response, we observed a trend toward higher response locked firing rates during conflict, but this difference did not survive cluster-wise correction for multiple comparisons ($p=0.07$, permutation test; Figure 5.5B, right).

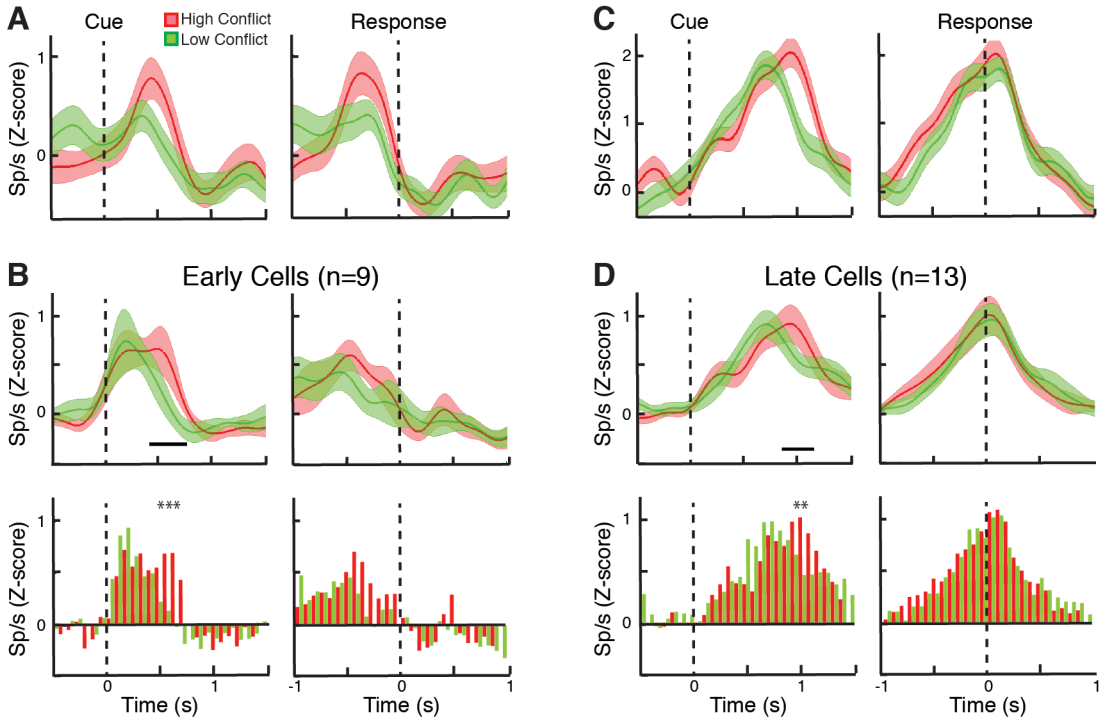


Figure 5.5. Conflict related differences in increased spiking activity in STN cells.
(A) Average across-trial firing rates (mean ± SEM) for correct low and high conflict trials for a single representative early cell, locked to cue presentation (left, t = 0) and response (right, t=0). **(B)** Average across-cell continuous firing rate (top, mean ± SEM) and spike histograms (bottom) for correct low and high conflict trials for all early cells (n = 9).

Time points and bins exhibiting a significant difference ($p < 0.05$ after correcting for multiple comparisons, permutation test) between low and high conflict trials are denoted by a black horizontal bar and by black *. Spike histograms are binned into 100 ms bins with 25 ms overlap. These cells demonstrated a delayed post peak decay of firing activity during high conflict trials. **(C)** Average across-trial firing rates (mean ± SEM) for correct low and high conflict trials for a single representative late cell locked to the cue (left, t = 0) and to the response (right, t = 0). **(D)** Average across-cell continuous firing rate (mean ± SEM, top) and spike histograms (bottom) for correct low and high conflict trials for all late cells (n = 13) locked to the cue (left, t = 0) and to the response (right, t = 0). During conflict, the rate of rise for the late cells was delayed.

Spiking activity for a representative late cell demonstrated an increase in firing rate that peaked with the motor response during both high and low conflict trials (Figure 5.5C). When time-locked to cue presentation, high conflict trials exhibited a delayed rise in spiking activity compared to low conflict trials, consistent with the delayed response time observed during the high conflict trials. Indeed, when the data were locked to the motor response we observed that the activity during both trial types peaked at their corresponding response times. When we examined all late cells, we found that this pattern of activity was consistent across this population of cells. Both continuous time firing rates and binned spiking activity peaked earlier during low conflict trials compared to high conflict trials when time-locked to cue presentation ($p < 0.05$, permutation test; Figure 5.5D), but when locked to the motor response, both trial types exhibited peaks at the same time. Though this suggests that the conflict related firing rate differences in these cells are merely due to differences in the timing of movement, there are a few features that suggest the possibility that other factors may also be at play. Across cells, we found that both low and high conflict trials exhibited spiking onset around 125 ms following cue presentation (difference in onset time $p > 0.05$, permutation test). As both trial types began firing at the same time, the delayed peak in spiking activity observed during high conflict trials may be related to rate at which firing rates increase in both trial types. To investigate this, we calculated the first derivative of the continuous time firing rate during each trial. We found significant differences in the rate of firing rate increases between low and high conflict trials that accounted for the observed differences in peak firing times ($p < 0.05$, permutation test; Figure 5.6).

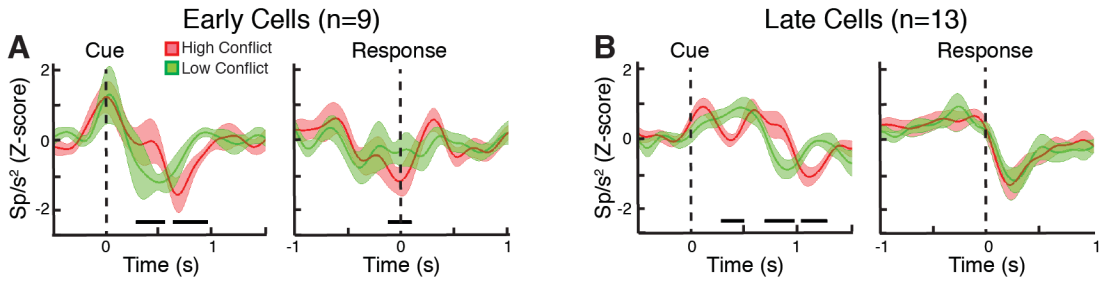


Figure 5.6. Conflict related differences in rate of increase in spiking activity in STN cells.
Same as **Figure 5.5**B,D, but the first derivative (d(sp/s)/dt) of the firing rate was calculated prior to averaging across trials and then across cells. Time points exhibiting a significant difference ($p < 0.05$ after correcting for multiple comparisons, permutation test) between low and high conflict trials are denoted by black horizontal bar.

We also examined the spiking response from all cells that exhibited a significant decrease in spiking activity during the decision period (n = 7; down cells). Representative spiking activity from a single down cell exhibited a decrease with the motor response (Figure 5.7A). Down cells demonstrated, on average, a decrease in firing immediately prior to the motor response (Figure 5.7B). We examined whether spiking activity in this population of down cells was modulated by decision conflict. Spiking activity for a representative down cell demonstrated an initial decrease during both high and low conflict trials that was maintained for a longer period of time during high conflict trials (Figure 5.7C). We found that across the population of down cells, this distinction between high and low conflict trials was also consistent and significant ($p < 0.05$, permutation test; Figure 5.7D, top). When we collected spiking activity in 100 ms bins, we found that spiking activity during high conflict trials was significantly less than during low conflict trials from 425 to 650 ms following cue presentation ($p < 0.05$, permutation test; Figure 5.7D, bottom).

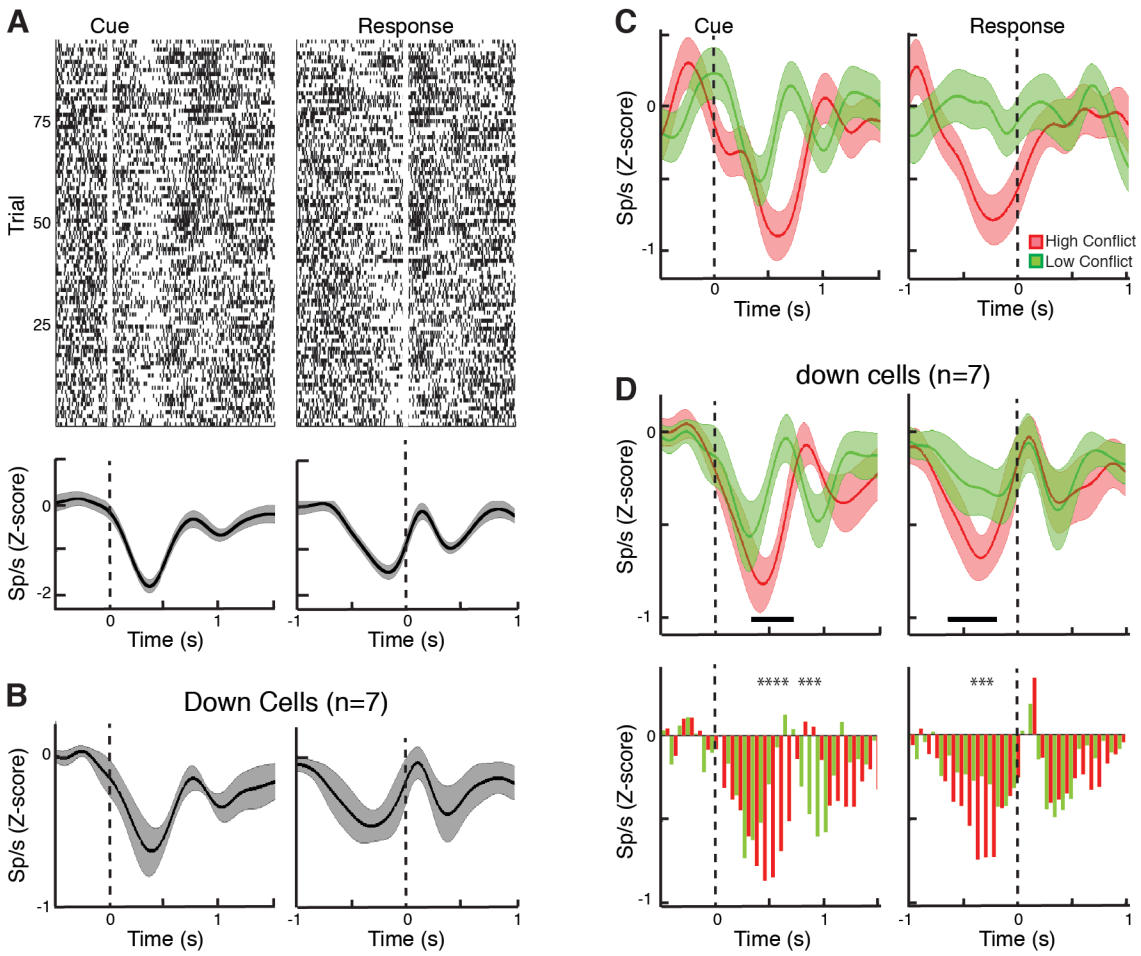


Figure 5.7. Conflict related differences in downward firing STN cells.
(A) A representative cell exhibiting decreased spiking activity following presentation of the target arrow, locked to the cue (left, t=0) and to the response (right, t=0). Raster plot (top) and the average continuous time firing rate (bottom, mean ± SEM) are shown for the cell. **(B)** Average continuous time firing rate (mean ± SEM) for all cells exhibiting a significant decrease in spiking activity (n = 7) locked to the cue (left) and to the response (right). **(C)** Average across trial firing rates (mean ± SEM) for correct low and high conflict trials for a single representative cell exhibiting a significant decrease in spiking activity locked to the cue (left, t = 0) and to the response (right, t=0). **(D)** Average across cell continuous firing rate (mean ± SEM, top) and spike histograms (bottom) for correct low and high conflict trials for all down cells (n = 7). Time points and bins exhibiting a significant difference ($p < 0.05$ after correcting for multiple comparisons, permutation test) between low and high conflict trials are denoted by a black horizontal bar and by black *. Spike histograms are binned into 100 ms bins with 25 ms overlap.

5.3.4 LFP-single unit interactions during conflict

In order to explore whether the observed changes in oscillatory LFP activity were linked to the changes in firing rates, we investigated the interaction between the oscillatory phase of the macroelectrode LFP signal and the microelectrode spiking activity. In a single cell, we observed that spiking activity occurred more frequently during the peaks of a theta oscillation (5-7 Hz; Figure 5.8A, top, left). When we computed the continuous firing rate during all high conflict

trials for this cell and time-locked this activity to the trough of every theta oscillation, we found that spiking activity was modulated by the phase of the theta oscillation (Figure 5.8A, middle, left). To quantify this relation, we tabulated the instantaneous theta phase for all spiking events that occurred during high conflict trials, and found that the distribution of phases was significantly non-uniform (R_{high} = 2.95, p = 0.002; see section 5.2.6 "Spike-phase interactions"; Figure 5.8A, right). In a separate cell, we observed that spiking activity occurred near the peaks of a beta oscillation (15-30 Hz; Figure 5.8B, top, left). When we time-locked the cell's continuous-time spiking activity during all high conflict trials to the trough of every beta oscillation, we found that the average firing rate for this cell was modulated by the phase of the beta oscillation (Figure 5.8B, middle, left). As in the theta locked cell, the distribution of instantaneous phases of the beta oscillation captured at each spiking event in the beta-locked cell was significantly non-uniform (R_{high} = 3.74, p < 0.001; Figure 5.8B, right).

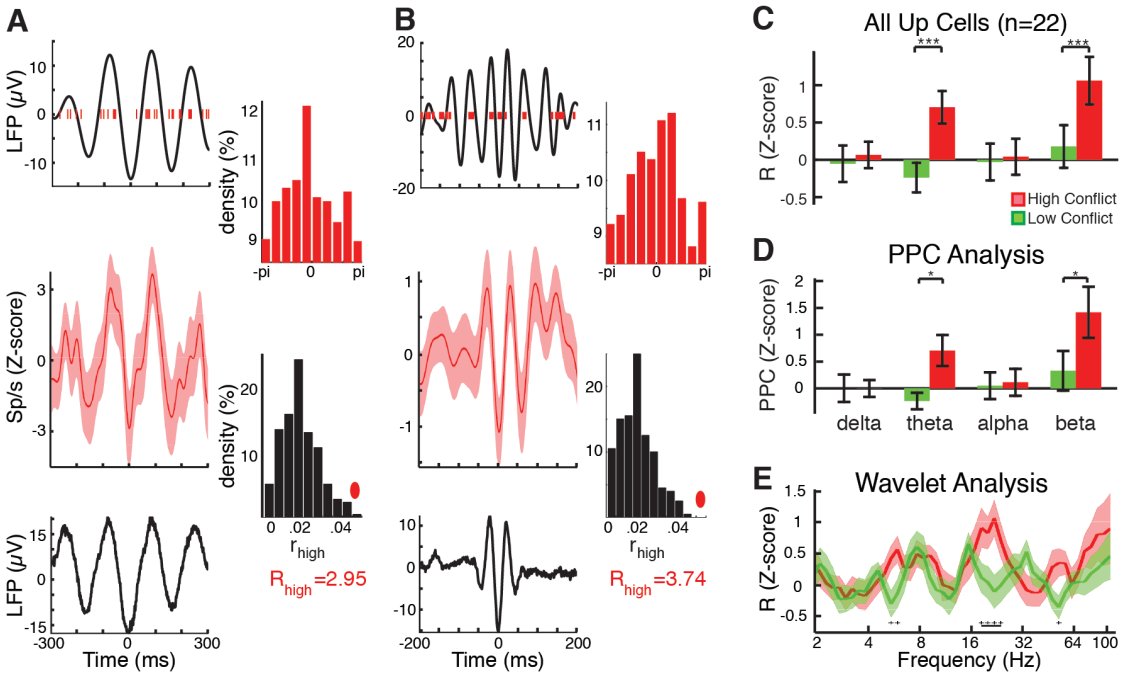


Figure 5.8. Theta and beta oscillations entrain STN cells during conflict.
(A) Spiking activity (red tick marks) in a single conflict triggered theta locked cell occurs more frequently during the peaks of a theta oscillation (top). Three cycles of a theta filtered (5-7 Hz) LFP trace are aligned to the trough of a theta cycle (black). Average firing rate (mean ± SEM), locked to all troughs of theta oscillations during high conflict trials, demonstrate significant modulations with theta phase (middle). The average raw LFP locked to the trough of theta oscillations is shown for reference (bottom). The histogram of instantaneous theta phases of each spiking event during

high conflict trials is non-uniform (top right). The theta spike phase locking value for this cell during high conflict trials (r_{high}, red dot) is shown in relation to 1000 permuted spike phase locking values calculated for this cell. The normalized spike phase locking value (R_{high}) demonstrates significant phase locking. **(B)** Spiking activity and phase locking for a single conflict triggered beta locked cell. Format is as in (A). **(C)** Across all cells that exhibited a significant increase in spiking activity (n=22) during the decision period, theta and beta spike phase locking were significantly greater during high compared to low conflict trials (mean ± SEM; *** denotes $p < 0.01$). (**D**) Same as (C) but with values derived using the pairwise phase consistency metric (mean ± SEM; * denotes $p < 0.05$). (**E**) Rzscore values derived from the instantaneous wavelet phase also revealed conflict related increases in spike phase locking in the theta and beta bands (mean ± SEM; + denotes significantly different locking between low and high conflict trials, $p < 0.05$; horizontal bar denotes contiguous frequency ranges that survived cluster-based correction for multiple comparisons)**.**

To assess whether there was a consistent pattern of spike phase locking across our population of cells, we calculated spike-phase locking for delta (2 - 4 Hz), theta (5 - 7 Hz), alpha (8 - 14 Hz), and beta (15 - 30 Hz) oscillations during high and low conflict trials for all 22 cells that demonstrated an increase in spiking activity during the task. We found that across these cells, the extent of spike phase locking during the high conflict trials was significantly greater than during the low conflict trials in the theta ($t(21) = 2.91$, $p = 0.008$) and beta $t(21) = 3.38$, $p = 0.003$) bands (Figure 5.8C). These differences arose because the theta and beta bands showed significantly non-zero spike phase locking during the high conflict trials (theta: $t(21) = 3.23$, $p = 0.004$; beta: $t(21) = 3.33$, $p = 0.003$) but not during the low conflict trials (theta: $t(21) = -1.21$, $p = 0.24$; beta: $t(21) = 0.62$, $p = 0.54$). Similar results were obtained if spike phase locking was quantified using pairwise phase consistency (Figure 5.8D) or the wavelet transform (Figure 5.8E). In order to determine whether the observed spike-phase interactions were time-locked to stimulus presentation or to the response, we examined the time evolving spike-phase locking in the theta and beta bands locked to the cue and to the response (Figure 5.9). We found significantly greater spike-phase locking in both frequency bands during high conflict trials compared to low conflict trials that was specific to the stimulus onset ($p < 0.05$, permutation test). Though the theta and beta band differences were significantly higher during conflict ($p<0.05$, permutation test) for both the cue and response aligned data, the response-aligned differences did not survive correction for multiple comparisons. We did not observe any significant spike phase locking with any other frequency band. We also did not find evidence for significant spike-phase locking in the down

cells ($p > 0.05$), suggesting that spike-phase locking was specific to conflict in the cells exhibiting an increase in spiking activity during the pre-response period.

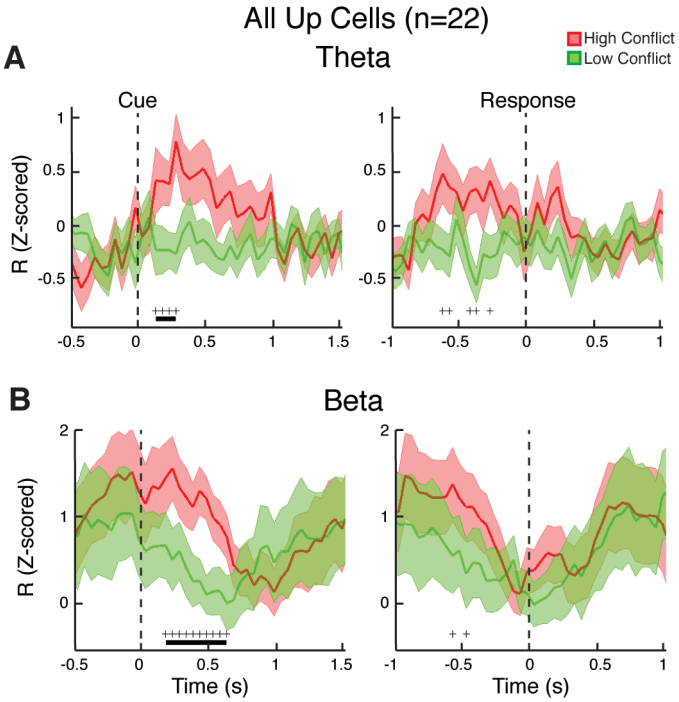


Figure 5.9. Time evolving theta and beta spike phase locking in STN cells.
(A) Across all cells that exhibited a significant increase in spiking activity (n=22) during the decision period, cue-locked (left, t=0) and response-locked (right, t=0) time evolving spike phase locking in the theta band was significantly greater during conflict (mean ± SEM; + denotes significantly different locking between low and high conflict trials, $p < 0.05$; horizontal bar denotes contiguous time points that survived cluster-based correction for multiple comparisons). (**B**) Same as (A) but for beta band. Beta spike phase locking was also significantly higher during conflict.

We subsequently analysed spike phase locking separately for early and late cells (Figure 5.10) in order to test whether the spike phase locking was specific to either of the two identified populations of cells. Early cells exhibited significant phase locking during the high conflict trials to theta and beta oscillations (theta $t(8) = 2.74$, $p = 0.026$; beta $t(8) = 2.51$, $p = 0.037$). Similarly, late cells also exhibited significant phase locking during high conflict trials to both theta and beta oscillations (theta $t(12) = 2.21$, $p = 0.047$; beta $t(12) = 2.28$, $p = 0.042$). Like the larger population, neither early nor late cells demonstrated significant theta or beta phase locking during low conflict trials ($p > 0.05$). Though the conflict related differences in the early and late cells were only significant in the theta and beta bands ($p < 0.05$, t-test; Figure 5.10A,C), respectively,

an ANOVA between cell type and conflict revealed a main effect of conflict but failed to reveal an interaction between conflict and cell type in either the theta (conflict $F(1, 20) = 8.26$, $p = 0.009$; cell type $F(1, 20) = 0.76$, $p = 0.39$; interaction $F(1, 20) = 0.14$, $p = 0.72$) or beta (conflict $F(1, 20) = 9.87$, $p = 0.005$; cell type $F(1, 20) = 0.20$, $p = 0.66$; interaction $F(1, 20) = 0.97$, $p = 0.34$) bands. When we analyzed the time evolving differences in early and late cells, we found that both populations of cells showed some time points that had significantly ($p<0.05$, permutation test) higher spike phase locking in the theta band, although none of these differences survived correction for multiple comparisons (Figure 5.10B,C). The beta band, on the other hand, showed a clear, significant increase during conflict for the late firing cells when locked to the cue (Figure 5.10D).

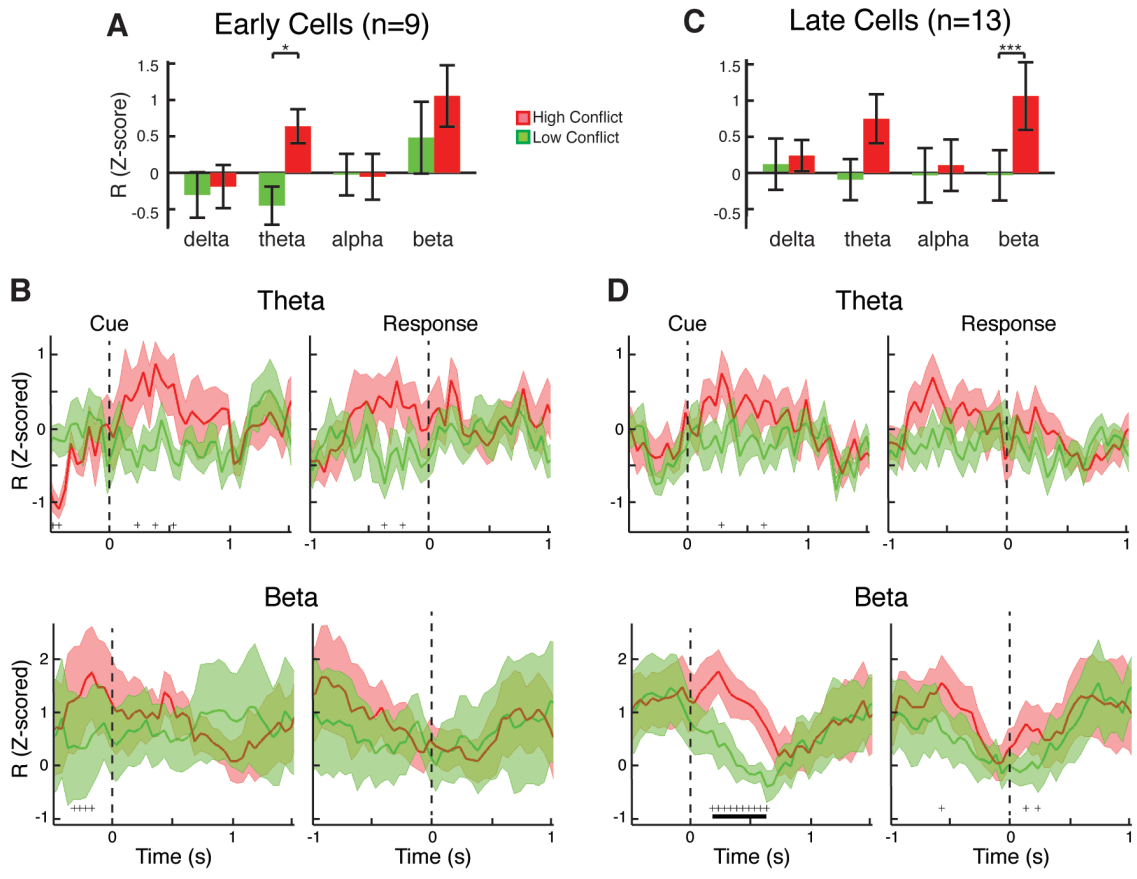


Figure 5.10. Early and late cell differences in spike phase locking.
(**A**) Same as **Figure 5.8**C but for just the early cells. (n=9, mean ± SEM). (**B**) Same as **Figure 5.9** but for just the early cells (mean ± SEM). (**C-D**) Same as (A-B) but for just the late cells (n=13). * denotes $p < 0.05$; *** denotes $p < 0.01$; +

denotes significantly different locking between low and high conflict trials, $p < 0.05$; horizontal bar denotes contiguous time points that survived cluster-based correction for multiple comparisons.

5.4 Discussion

Consistent with previous studies (Cavanagh et al., 2011; Fumagalli et al., 2011; Brittain et al., 2012; Zaghloul et al., 2012; Zavala et al., 2013), we found increased theta oscillatory power and related changes in spiking activity in the STN during decisions that involved high conflict. Importantly, by directly measuring simultaneous LFP and spiking activity in the human STN, our data demonstrate significant interactions between changes in theta and beta oscillations and spiking activity during these difficult decisions. We also showed that when these interactions were disrupted by DBS, more errors occurred during high conflict trials. Though it is important to note that the ITPC differences reported here were inconsistent with the ones reported previously in Chapter 3, we will address this issue further in Chapter 6 below.

We and others have shown consistent theta band increases in the STN during conflict. It was previously unknown, however, how these oscillations interacted with the firing rate changes of the STN. Here we demonstrate the relation between theta oscillations and individual firing rates of STN neurons during the decision process. Across all neurons that showed task related increases in firing rate, theta oscillations significantly entrained spiking activity, but only during conflict. We identified a population of STN neurons that exhibit an early increase in activity during all trials and that maintain their elevated firing rates for a longer period of time during trials that involved high conflict. Computational models have predicted that early increases in STN spiking activity during a decision may represent an anti-kinetic signal, conveyed directly from the cortex to the STN, that initially halts all motor responses at the beginning of each decision (Nambu et al., 2002; Frank, 2006). These models have further predicted that during conflict, this early firing rate increase would persist for longer to inhibit movement and thus "hold your horses". From this perspective, our data suggest that greater theta oscillatory entrainment

observed in the early cells during high conflict trials may be important in maintaining elevated STN firing rates during difficult decisions, thereby prolonging the inhibition of movements.

We also identified a population of STN neurons that exhibited increases in spiking activity that peaked during the motor response, as well as a population of neurons that decreased spiking activity following the cue. As STN firing is presumed to inhibit movement (Albin et al., 1989; DeLong, 1990), the decreases in firing we observed in the latter population would be consistent with the release of movements associated with the selected response. However, the observed movement related increases in spiking activity in some STN neurons, as well as in their downstream targets (Goldberg et al., 2013), suggest a more complex organization. One possibility is that basal ganglia circuits are organized in multiple parallel motor loops with a center-surround architecture (Nambu et al., 2002). In this paradigm, movement related increases in late cell firing may reflect the inhibition of actions that might interfere with the execution of the appropriate response. Consistent with such a center-surround organization, we found a subset of neurons that exhibited significant differences in spiking activity depending on the direction of movement. Another possibility, however, is that the movement related increases in STN firing may instead convey a pro-kinetic signal that could inhibit the GPi via the GPe (Isoda and Hikosaka, 2008). Regardless of whether the late cells reflect pro- or anti-kinetic signalling, the activity of this population of late cells is clearly related to the motor response. Notably, we found that the level of conflict determined the rate at which their spiking activity changed. Because we did not capture eye movement, hand accelerometer or EMG activity, however, our data do not establish whether these changes in firing rate increases are a cause or a consequence of delayed eye or hand movement. Nevertheless, it is worth noting that similar conflict-related changes in firing rate have been observed in cortical structures during perceptual action selection (Gold and Shadlen, 2002, 2007). Thus our data raise the possibility that the process of cortical neuronal

evidence integration governing motor responses during decisions may also be reflected in the spiking activity of these STN neurons.

As most of our previous work focused on changes in STN theta activity, our finding that STN neurons also locked to beta oscillations during conflict was slightly unexpected. This finding, however, is strongly in line with the anti-kinetic role that has been proposed for beta oscillations (see section 1.1.5 “Pathological oscillations in the basal ganglia” and section 1.4.2 “Synchronized oscillations in the beta frequency (13-30 Hz) band”). All of the tasks presented in this thesis consistently showed a decrease in beta band activity prior to movement, which is in agreement with the STN literature. Furthermore, increases in beta band activity in the STN (Kühn et al., 2004; Leventhal et al., 2012; Ray et al., 2012; Alegre et al., 2013) have been observed during response inhibition tasks in which movements must be cancelled. Finally, one study has shown that high conflict trials of the Stroop task are associated with elevated beta levels just before a response. In that report, both low and high conflict trials show a large decrease in beta when the stimulus was shown (see Figure 1.2), but during conflict the beta power decrease slightly reversed prior to the response. During error trials, the small beta power reversal also occurred, but it took place after the error was committed. The authors interpreted these findings as suggestive that there are subcircuits within the motor network that can enhance beta power to inhibit response even though the overall network is showing a decrease in beta synchronization. In line with this idea, we found decreased macroelectrode LFP beta power coincident with a conflict-triggered increase in beta band spike phase locking. This suggests that although spatially extensive beta synchronization generally drops, there are some neurons between which beta synchrony actually increases during conflict to delay movement. One hypothetical mechanism by which these beta increases could inhibit movement is by entraining the movement related cells discussed in the previous paragraph and delaying the rate at which they rise to threshold. This claim is supported by previous studies showing that focal drops in striatal beta activity allow cells in that area to

increase their firing rate in a movement related manor (Courtemanche et al., 2003) and that beta power and neuronal firing are inversely related (Spinks et al., 2008). Thus, during conflict, focal increases in STN beta activity may re-engage movement related STN cells to prevent the rate coding of an action.

The source of such focal conflict related beta band synchronization, unlike that in the theta band, has not been clearly established. One potential source is input from the inferior frontal cortex (IFC). The IFC projects to the STN (Aron et al., 2007), both the STN and IFC show similar beta power increases during successful stopping (Swann et al., 2009, 2011, 2012; Ray et al., 2012), and connections between the STN and IFC are enhanced during response inhibition (Aron et al., 2007; Jahfari et al., 2012; Herz et al., 2014). Indeed, previous imaging studies and computational models have suggested that the mPFC and IFC activate the STN in parallel, rather than in sequence, during decision-making (Wiecki and Frank, 2013; Herz et al., 2014). As theta and beta oscillations also mediate communication between the mPFC and IFC (Hanslmayr et al., 2007; Cavanagh et al., 2009; Cohen and Cavanagh, 2011; Ruiz et al., 2011; Swann et al., 2012), our data suggest that the interactions we observe in the STN may subserve a larger cortical-subcortical network activated during decision conflict (Aron et al., 2007; Neubert et al., 2010; Herz et al., 2014).

Taken together, our results suggest that theta and beta oscillatory activity interact with enhanced neuronal spiking to delay movement through two distinct mechanisms. The first involves increases in theta oscillatory power and a related entrainment of neurons during high conflict trials. These changes are associated with a more persistent elevation of discharge rates amongst early cells, thereby prolonging motor inhibition and buying time for the selection of the most appropriate behavioural response. Parallel to this, a second process involves beta band suppression, focal increases in beta entrainment during high conflict trials, and an increase in late cell discharge rates that may be linked to the inhibition of conflicting response alternatives or to

executing the desired response through GPe mediated inhibition of GPi. These data therefore suggest the presence of STN neuronal entrainment to two separate frequency bands during conflict, which may facilitate specific channels of communication with different cortical and subcortical structures during the decision process (Fries, 2005). Consistent with this architecture and with previous reports (Frank et al., 2007; Cavanagh et al., 2011), when we disrupted these networks post operatively with STN DBS, participants exhibited more errors in this simple decision task.

Chapter 6. Human subthalamic nucleus-medial frontal theta phase coherence is involved in conflict and error related cortical monitoring

6.1 Introduction

The work described in the previous chapters showed that during low and high conflict action selection the STN theta phases realign across trials when a cue is shown and STN theta power increases. During conflict, mPFC and STN phases become coherent and the theta phases of STN LFP oscillations entrain neuronal populations. These findings all suggest that theta activity plays an important role in mediating conflict related information between brain areas, but there are a few remaining questions that our previous studies were unable to address. Though the design of the dots task (Chapter 4), with its gradual onsets and offsets of conflict, allowed us to show that STN theta activity rises and falls during action selection as conflict appears and disappears, many events in everyday life require an immediate decision in the face of conflict (i.e. Should I run to catch the bus or should I just wait for the next one) and do not give us the luxury of gradually showing us their conflicting nature and then gradually showing us the correct answer. Thus it remains to be determined if the increased synchrony between the mPFC and the STN we reported in Chapter 4 also occurs during rapid onsets of conflict. Related to this question, it also remains unknown what relationships may or may not exist between across-trial theta phase resets in one brain structure and across-site theta phase coherence between different brain structures. The dots paradigm involved very gradual changes in the degree of conflict, not rapid onsets as in the flanker task. Because of these gradual changes, the dots paradigm did not involve a "cue" or stimulus onset, thus there was no cue-related theta phase reorganization as indexed by ITPC. As both flanker studies reported above did show task related increases in ITPC, it remains to be determined whether or not these increases would interfere with or enhance the inter-site phase coherence, which we showed can occur without ITPC increases (Chapter 4). Understanding the

interactions between these two phenomena may also help explain the contradictory findings reported in the previous two chapters related to ITPC: one study showed lower ITPC values for the slowest high conflict trials, while another study showed higher ITPC values for the high conflict trials relative to the low conflict trials.

A final outstanding question that we sought to answer in this study centers around the important role the mPFC seems to play in across-trial adaptations to the level of conflict or to errors (Kerns et al., 2004; Ridderinkhof et al., 2004; Danielmeier and Ullsperger, 2011). mPFC theta activity seems to interact with areas of the dorsolateral prefrontal cortex (DLPFC) following high conflict and error trials, which may be related to the respective speeding and slowing of reaction times on subsequent trials (Kerns et al., 2004; Hanslmayr et al., 2007; Cavanagh et al., 2009; van Driel et al., 2012). Whether or not these interactions also involve the STN remains an open question. A link is suggested, however, by studies showing error related activity in the STN (Brown et al., 2006; Zavala et al., 2013; Cavanagh et al., 2014; Siegert et al., 2014). Recently Cavanagh et al. demonstrated that STN theta phases recorded during the Simon task are related to the post error slowing that occurs on a trial that follows an error. Given that we previously demonstrated that theta oscillations can connect the mPFC and STN during conflict, we hypothesized that a similar coherence may be at play during trials that follow an error and thus are associated with more cognitive control. Here, we test this issue, as well as those raised in the previous paragraph, by simultaneously recording electrophysiological activity from the mPFC and STN while a fourth group of DBS patients once again performed the Eriksen flanker task.

6.2 Methods

6.2.1 Subjects and task

All subjects gave their written informed consent to take part in the study, which was approved by the appropriate local ethics committees. Thirteen subjects (13 males; mean disease duration, 10

years; mean age, 56 years; age range, 42–69 years) underwent bilateral implantation of DBS electrodes into the STN, as a prelude to high-frequency stimulation for the treatment of advanced PD. Clinical details of the patients are available in Table 6.1. The mean percentage improvement in the motor section of the Unified Parkinson's Disease Rating Scale (UPDRS) following treatment with levodopa was 64.4 ± 4.5% ($p < 0.001$, Wilcoxon signed-rank test between ON and OFF levodopa scores) across subjects, indicating good responsiveness to levodopa in our study participants. Patients performed an arrow version of the flanker task (Eriksen and Eriksen, 1974) while receiving their regular medication 3–6 d after electrode implantation. The task was identical to one described in Chapter 3 above (Zavala et al., 2013), a schematic of which is shown in Figure 6.1A. Data from two of the participants (cases 1 and 2) were also included in our previously published work presented in Chapter 3 (Zavala et al., 2013), although that work did not include mPFC theta connectivity, which is the main focus of this chapter.

Table 6.1. Clinical details from third Flanker experiment

Case	Age	Disease Duration	UPDRS Off (III)	UPDRS On (III)	First Symptom	Reasons for Surgery	Daily Medication (mg/d)
1	58	10	42	20	Leg cramp	Tremor	Trihexyphenidyl 3; Levodopa 600; Rasagiline 1; Amantadine 100
2	62	10	20	8	Left side tremor & bradykinesia	On/off fluctuations, tremor & impulse control disorder	Levodopa 1000; Trihexyphenidyl 6
3	61	4	37	15	Left side tremor	Tremor	Amantadine 200; Levodopa 750; Entacapone 1000
4	65	15	51	21	Left hand tremor	Freezing	Amantadine 200; Levodopa 400; Ropinirole 12
5	44	10	33	7	Left hand tremor	Motor fluctuations, dyskinesia	Amantadine 200; Ropinirole 24; Rasagiline 1; Levodopa 600; Apomorphine 4.5mg/hr
6	42	9	60	42	Loss of dexterity	Bradykinesia, dystonia, freezing	Amantadine 400; Levodopa 600
7	51	9	56	12	Tremor	Motor fluctuations, dyskinesia	Apomorphine 4; Levodopa 1500; Entacapone 600

8	43	10	36	6		Motor fluctuations	Levodopa 600; Ropinirole 2
9	59	16	54	9	Loss of dexterity	Dyskinesia, painful cramps	Levodopa 100; Apomorphine 5mg/hr Rotigotine 4
10	69	10	34	12	Shuffling gait	Motor fluctuations, bradykinesia	Levodopa 800; Amantadine 100
11	53	7	25	5	Loss of dexterity	Dyskinesia, bradykinesia	Levodopa 800; Entacapone 800; Rasagiline 1
12	67	12	25	9		Off periods & dyskinesia	Levodopa 950; Amantadine 100; Pramipexole 2.25
13	63	13	32	18	Dragging of left leg	Stiffness	Levodopa 700; Entacapone 1000; Ropinirole 8; Quetiapine 25; Clonazepam 0.5

6.2.2 LFP data recording and analysis of power

STN LFPs were recorded from the DBS electrodes. Simultaneously, continuous scalp EEG was recorded from frontal, central and parietal electrodes over the midline (Fz, Cz and Pz; International 10-20 System). More lateral electrodes were prohibited by surgical wounds and dressings in this patient group. All signals were sampled at 2048 Hz, bandpass filtered between 0.5 and 500 Hz, and amplified using a TMSi Porti and its respective software (TMS International, Netherlands). Monopolar recordings were subsequently converted off-line to a bipolar montage between adjacent contacts (three bipolar channels per STN side and two bipolar channels for the EEG recordings: Fz-Cz and Pz-Cz) to limit the effects of volume conduction from distant sources. Before further analysis, LFP data were filtered between 1 and 500 Hz and down-sampled to 1000 Hz.

Any trials with reaction times >1.5 s (including no response trials) or <300 ms were not included in the analysis (<1%). For the comparison of correct low conflict and high conflict trials, all incorrect responses were excluded. For the comparison of high conflict error trials to high conflict correct trials, a reaction time matched subset of the correct trials was used due to the very low number of errors trials we observed. For each incorrect high conflict trial, one reaction time

matched trial was randomly chosen from all of the correct high conflict trials that had a reaction time within 30 ms of the incorrect trial. Any incorrect trial that did not have a correct trial that was within 30 ms was excluded from the analysis. Two subjects had fewer than 5 errors and were therefore not included in this analysis. The average number of error trials for the 9 subjects that were included in the error analysis was 8 ± 0.8 trials. For the analysis shown in Figure 6.3, the correct trials that followed error trials were compared to the correct trials that followed the reaction time matched set of correct trials. The average number of trials used for this analysis was 6 ± 0.9 trials for the post error trials and 6.1± 0.8 for the post reaction time matched correct trials. These numbers were slightly lower than the number of errors trials and reaction time matched correct trials because a few subjects had consecutive errors or errors that occurred at the end of a block. The post error trials and the post reaction time matched correct trials had a similar proportion of high conflict trials that was consistent with the 2:1 ratio of high to low conflict trials that was used in the task (63.5± 8.1% and 64.3± 4.3% for the post reaction time matched correct trials, $p=0.85$, paired t-test).

Two of the subjects showed significantly higher error rates than the other 11 subjects (error rate across all trials=27.7% and 27.8% for the two outlier subjects and 8.5±1.1% for the 11 other subjects, $p<0.001$, unpaired t-test) and were therefore excluded from most of the analysis. Prior to the exclusion of these two subjects, however, permutation testing was used to test if the behavioural differences between low and high conflict trials were significant when all subjects were included. Furthermore, a 2 (hemisphere) x 2 (conflict level) repeated measures ANOVA was also used to test if there was a hemisphere-dependent effect of conflict on theta (4-8 Hz) and beta (14-32 Hz) power when the outlier subjects were included in the analysis. This was done by extracting the time-evolving wavelet power for each trial using the Morlet wavelet at 8 scales/octave from 2 to 107 Hz. A time window of interest (1 s long beginning at arrow onset) was then selected, and the average theta and beta power during that window was found separately

for the low and high conflict trials in each subject's left and right STN. One subject's recordings were only accessible from one STN; therefore that subject's data were not included in the repeated measures ANOVA.

For all of the subjects that were not excluded due to poor behavioural performance (n=11), the instantaneous theta power and phase of the bipolar LFP and EEG signals were calculated by bandpass filtering each trial's raw signal between 4 and 8 Hz and applying the Hilbert transform. Each trial was analysed from 0.75 s before to 2 s after arrow onset for the cue-aligned analysis, from 1.5 s before to 1.5 s after the response for the response-aligned analysis, and from 2 s before to 0 s before arrow onset for the analysis of the pre-stimulus period (Figure 6.4). A 1 s buffer on either side was used when calculating phase and power to eliminate any edge effects. Any trial with a clear artefact in any of the LFP or EEG channels was discarded.

To assess differences in theta power between low and high conflict trials, the following approach was used. First, the mean power in each bipolar recording for each trial type was calculated by averaging the theta power time series across trials. This method produced a time series for low and high conflict trials for each of the three bipolar contacts on each STN electrode and the two bipolar EEG contacts. Each time series was then normalized to the mean power of that channel recorded during a "baseline" period of all trials in the full second leading up to the warning cue onset (t=-1.5 to -0.5, relative to the arrow onset). Finally, all three STN bipolar contacts were averaged together, and the difference between the low and high conflict trials was calculated (see figure Figure 2.1). Averaging across all the contact pairs in a given electrode was performed so as to avoid selection bias, although this procedure might serve to underestimate spectral changes in the STN as not all contact pairs were necessarily within or bridging this nucleus. The above procedure was repeated for the electrode in each STN to find the mean difference across subjects. To assess the statistical significance of this difference, the same procedure was repeated 1000 times with the low and high conflict trial labels randomly assigned during each permutation. The

p value of each time point was found by comparing the actual mean difference to the distribution of the 1000 permutations. The p values were then corrected for multiple comparisons using exceedence mass testing (Maris and Oostenveld, 2007), see methods section 2.3.2 “Permutations and exceedance mass testing” above. When performing other comparisons (i.e., high conflict errors vs. high conflict correct, post high conflict errors vs. post high conflict correct, fast-high conflict vs. slow-high conflict), the same procedure was repeated using the relevant trial groups.

To assess the single-trial correlation between reaction time and normalized power changes, the normalized theta power of each trial was averaged across the time period 1.25 ms before the warning cue onset and correlated with the reaction time in that trial. The resulting correlation coefficients (positive and negative coefficients derived using Spearman's correlation) were then averaged across the STN sides and a two-tailed, one-sample t-test was performed to determine whether the mean correlation was significantly different from zero across subjects.

6.2.3 Intertrial phase consistency (ITPC)

ITPC values were calculated separately for low and high conflict trials. In order to prevent the 2:1 high to low conflict trial ratio from affecting our results, the high conflict trials were down-sampled for each subject to match the number of low conflict trials. 1000 down-sampled ITPC time series were calculated for each subject’s high conflict trials, and the average of the down-sampled values was used for each subject. The low conflict and the down-sampled high conflict ITPC time series were then normalized to the “baseline” ITPC value recorded during the full second leading up to the presentation of the warning cue onset (t=-1.5 to -0.5, relative to the arrow onset). Randomly drawn subsets of all the trails were used to calculate the baseline 1000 times, with the number of trials used for each baseline calculation equal to the number of low conflict trials (and the number of down-sampled high conflict trials). An identical procedure was used when comparing error trials to correct trials, with the only exception being that the number of trials used to calculate the baseline was equal to the number of error trials.

To assess the statistical differences between conditions, the difference between the two normalized ITPCs was first calculated for each bipolar signal and averaged across all three bipolar contact pairs of the STN electrodes before being averaged across subjects. The same procedure was then repeated 1000 times with the labels of the two trial groups being randomly assigned at the beginning of each permutation. The p value at each point was calculated using the distribution of the 1000 permuted values and corrected for multiple comparisons at a significance level of 0.05 using exceedence mass testing.

6.2.4 Intersite phase coherence

The cortico:STN phase coherence was calculated using the continuous time evolving methods we described in the methods section 2.2.5 “Inter-site Lachaux coherence” above (Zavala et al., 2014) as outlined by Lachaux et al. (2002). Low and high conflict trials were analysed separately, with the high conflict trials phase coherence values being calculated 1000 times using a down-sampled data set. The difference between the instantaneous theta phase (projected on the complex plane) at time t in each bipolar STN contact and the theta phase at time t in the Fz-Cz channel was found at each time point. The phase difference values at time t were then averaged across trials and a sliding window was used to integrate across time (Lachaux et al., 2002). The width of the window was chosen to be 333 ms (2 cycles of a 6 Hz oscillation). The magnitude of the resulting average was then taken to generate the phase coherence. Each channel’s time evolving phase coherence signal was then normalized by that channel's “baseline” phase coherence. The baseline was chosen in the same way as it was chosen for the ITPC analysis: calculating the mean phase coherence value (averaged across the full second before the warning cue onset) 1000 times for down-sampled sets of trials (equal in number to the number of low conflict trials) randomly chosen from all of the trials. The three resulting normalized time series generated for each of the three contacts in each STN were then averaged within each STN before averaging across STNs. Statistical significance was determined using permutation testing as outlined above. To calculate

the Hilbert parietal cortico:STN phase coherence, the same analysis was done using the Pz-Cz bipolar electrode instead of the Fz-Cz electrode. An identical procedure was used when comparing error trials to correct trials, with the only exception being that the number of trials used to calculate the baseline was equal to the number of error trials.

6.3 Results

6.3.1 Behavioural Effects

As we recorded STN LFP activity and cortical EEG activity, subjects performed the exact same flanker paradigm that was used in our previous study described in Chapter 3 (Zavala et al., 2013). Consistent with our previous results, patients were significantly slower during the high conflict condition relative to the low conflict condition (mean ± SEM= 623±38 ms vs. 544±34 ms, $p<0.001$, permutation testing) and showed a significantly higher error rate during conflict (mean ± SEM = 14.5±2.5% vs. 5.2±2.8%, $p<0.01$, permutation testing; Figure 6.1A, inset). There was no slowing of reaction time in trials that followed an error ($p>0.05$, permutation testing). There was also no Gratton effect on reaction time (Gratton et al., 1992): the reaction time of high conflict trials was not affected by the level of conflict in the previous trial ($p>0.05$, permutation testing), and the same was true for low conflict trials (ANOVA, within-subject repeated measures, current trial conflict x previous trial conflict: current trial conflict $F = 32.0$, $p <0.001$; previous trial conflict $F = 0.3$, $p = 0.59$; interaction $F = 0.13$, $p = 0.72$).

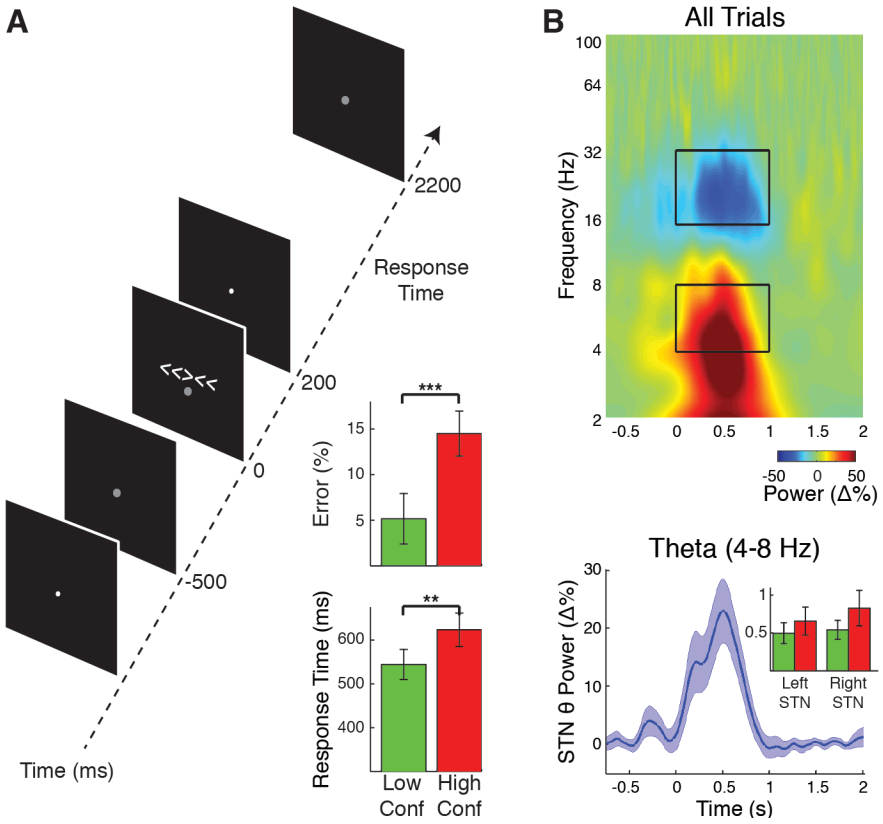


Figure 6.1. Task, behaviour, and theta.
(A) Patients performed an arrowed version of the Eriksen Flanker task, where they indicated the direction of a middle arrow flanker by either low (<<<<<) or high (>><>>) conflict arrows. A fixation dot was present throughout the entire experiment, but changed from white to grey 500ms before the arrows were shown. Inset shows that subjects had slower reaction times and higher error rates during the high conflict trials. Group average ± SEM is shown. ** denotes $p<0.01$, *** denotes $p<0.001$. **(B)** Consistent with previous results, the onset of the arrows elicited a decrease in beta power and an increase in low frequency, including theta, power. Our subsequent analysis focused on the band passed theta (4-8 Hz) power (bottom). Across subject group average is shown for top and bottom panels, with shaded region in bottom panel denoting SEM across subjects (n=13). Average theta power (inset) during the time window shown in the top panel revealed a significant effect of conflict that did not depend on hemisphere.

6.3.2 Conflict related difference in STN LFP and mesial frontal theta activity

In agreement with the results we reported in Chapter 3 (Zavala et al., 2013), subjects demonstrated an increase in theta power in the STN LFP during the task, as well as a decrease in beta power (top panel Figure 6.1B). As most of the literature concerning cortical and subcortical conflict related networks implicates coupling in the theta band (Cavanagh et al., 2011, 2012; Zavala et al., 2013, 2014), we focused our attention on this band. The bottom panel of Figure 6.1B shows that during the task, the subjects showed a consistent increase in band-passed theta (4 to 8 Hz) power in the STN LFP after onset of the warning cue (at t=-500 ms), and this was

followed by an even greater increase after onset of the arrows (at t=0 ms). A 2 (hemisphere: left vs. right) x 2 (conflict: high vs. low) repeated-measures ANOVA was used to confirm the results we presented in Chapter 3. When analysing the theta power during the window shown in the top panel of Figure 6.1B, we found a significant effect of conflict, but no effect of hemisphere and no hemisphere-by-conflict interaction (conflict $F = 5.72$, $df = 1$, $p = 0.036$; hemisphere $F = 0.41$, $df = 1$, $p = 0.53$; interaction $F = 0.91$, $df = 1$, $p = 0.36$). As there was also a decrease in beta power during the task, we also ran a repeated-measures ANOVA to test for differences in beta power, however, this revealed no significant effects of conflict or hemisphere (conflict $F = 0.93$, $df = 1$, $p = 0.36$; hemisphere $F = 0.04$, $df = 1$, $p = 0.84$; interaction $F = 0.27$, $df = 1$, $p = 0.61$). We therefore focused the remained of our analysis on just the theta band.

In order to explore the temporal evolution of theta band activity during the flanker task, we analysed the theta power, theta ITPC, and theta inter-site phase coherence in the STN LFP and in the mesial frontal EEG (Figure 6.2). Consistent with our previous results, we showed that high conflict trials were associated with a higher increase in pre-response theta power in the STN LFP relative to the low conflict trials (Figure 6.2A, top). There was also a trend toward higher post response STN theta power during the period that followed high conflict trials but before onset of the warning cue for the next trial (see section 6.3.4 “Pre-stimulus correlates of conflict behaviour“ below). When we analysed the mesial frontal theta power changes (Figure 6.2A, bottom), we observed no significant conflict related differences. Surprisingly, the mesial frontal theta activity peaked after the response (mean peak time ± SEM for all correct trials= 85.9±78.6 ms relative to the response), which was significantly later than the STN theta power peak that occurred before the response (mean peak time±SEM for all correct trials = -58.4±28.1 ms relative to the response, $p<0.05$, permutation testing; relative to the cue these values were 770 ± 82.5 ms for the mesial frontal EEG and 633.7 ± 94.1 ms for the STN LFP, $p<0.05$, permutation testing).

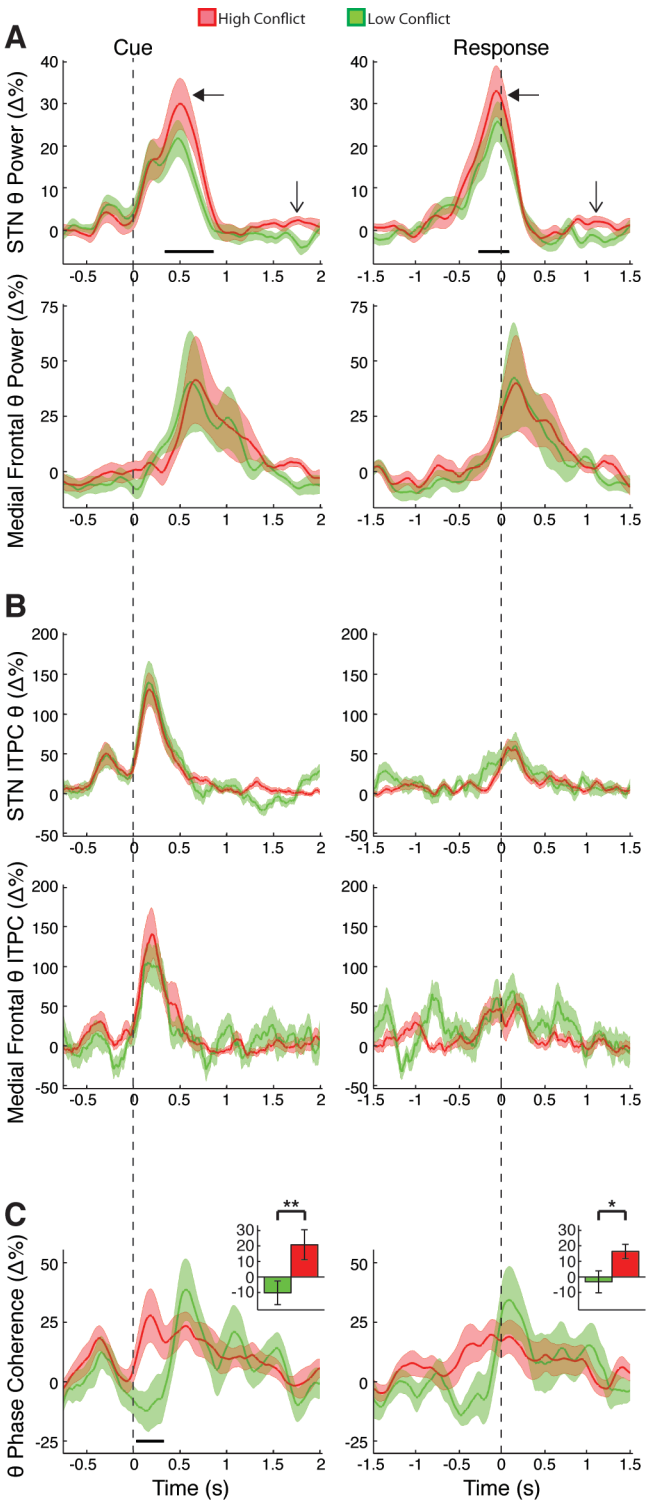


Figure 6.2. Conflict related theta activity in the STN and mPFC.

(A) Top row shows average STN (n=21) theta power aligned such that t=0 corresponds to the stimulus onset (left) or the response (right). High conflict trials showed significantly higher pre-response STN theta power (see solid arrows and black bars). A trend towards post response differences can also be seen as the post high conflict theta power is higher during the "baseline" period of the subsequent trial (see line arrows). Black line denotes significant time points ($p<0.05$ after correcting for multiple comparisons, permutation test). Group average percentage change ± SEM are

shown for all plots. Bottom row is same as top row but for mPFC theta power (n=11 recordings). Note STN theta power peaked before the response while mPFC theta power peaked after the response. **(B)** Same as A but for inter-trial phase consistency (ITPC). There were no significant time points that showed a conflict related difference. Note both the STN and mPFC increases in ITPC occurred at the same time and earlier than the theta power increases in either site. **(C)** Same as A but for the inter-site phase coherence. High conflict trials showed significantly higher pre-response coherence between the STN and the mPFC. This increase occurred at the same time as the increase in ITPC, suggesting that though both low and high conflict trials show an increase in phase consistency early in the trial (before the theta power increase), only during conflict are the phases synchronized between the two brain regions. Insets show the group average ± SEM phase coherence in the 300 ms following the cue (left) or the 500 ms preceding the response (right). * denotes $p<0.05$, **denotes $p<0.01$.

We have previously shown that the presentation of the warning cue or of the target arrows consistently realigns the phases of ongoing STN LFP theta oscillations as indexed by ITPC (Chapter 3 and Chapter 4 above). Here we reproduce these results and show that a simultaneous increase occurs in the mesial frontal EEG. Both STN LFP and mesial frontal EEG ITPC increases were much higher for the cue-aligned data than they were for the response aligned data, supporting our previous claim that the ITPC changes are evoked by stimulus onset. Notably, the STN LFP and mesial frontal EEG ITPC increases seemed to occur simultaneously, early during the trial (mean STN LFP ITPC peak time ± SEM for all correct trials = 257.6±49.8 ms relative to the cue; peak time for medial frontal EEG ITPC = 319.5 ± 115.6 ms, $p>0.58$, permutation testing), and they both peaked significantly ($p<0.001$, permutation testing) before the theta power increases described above. Though there were no significant, conflict related differences in ITPC in either location, this may be because it is not the particular phase of the oscillation that is important, but rather the coherence of the phases between the STN LFP and mesial frontal EEG.

Consistent with the idea that coherent oscillations between two brain sites would allow behaviourally relevant information to flow from one site to the other (Fries, 2005), we observed higher pre-response theta phase coherence between the mesial frontal EEG and the STN LFP during the high conflict trials (Figure 6.2C). Averaging across the first 300 ms following the cue showed a phase coherence value of 20.6±9.6% for the high conflict trials and a value of -10.1±7.6% for the low conflict trials ($p<0.01$, permutation testing). Likewise looking at the 500 ms leading up to the response showed higher phase coherence values for the high conflict

condition relative to the low conflict condition (16.4±4.6% vs. -3.1±7.0%, $p<0.05$, permutation testing). As a control, we also analysed the theta band phase coherence between the STN LFP and EEG over the parietal cortex (Pz-Cz) and found no conflict related differences (data not shown). The conflict related increase in phase coherence was not secondary to ITPC differences, as neither the STN LFP nor mesial frontal EEG showed any significant differences in ITPC for low or high conflict trials. However, given that the stimulus onset induced a phase reset in both the STN LFP and the mesial frontal EEG for low conflict trials, it is interesting that these trials did not show any increases in phase coherence during the periods with elevated ITPC levels. This discrepancy, together with the fact that the phase consistency across trials is not uniform (average maximum un-normalized ITPC across all correct trials in all subjects=0.34±0.03; perfect inter-trial phase alignments would have a value of 1) shows that there is some jitter in the exact phase to which both structures are aligning. Only during conflict did the two structures reset their phase in a way such that phase differences were sustained across time both within and across trials.

6.3.3 Post Error related difference in STN LFP and mesial frontal EEG theta phase coherence

Much of the literature concerning the mPFC focuses on the role this brain area might play in error monitoring and post error adaptations (Ridderinkhof et al., 2004; Danielmeier and Ullsperger, 2011). When we compared the theta band activity in the incorrect high conflict trials to the activity observed in a reaction time matched set of correct high conflict trials, we observed no differences in theta band power, ITPC, or phase coherence (data not shown). We also observed no differences in theta power or ITPC when we compared the correct trials that followed a high conflict incorrect trial to the correct trials that followed a reaction time matched high conflict correct trial (Figure 6.3A,B). However, when we analysed the phase coherence between the mesial frontal EEG and the STN LFP, we observed a significant, pre-response difference between the trials that followed errors and those that did not. Averaging across the same windows that

showed conflict related differences in phase coherence (t=0 to 300ms for cue-aligned data or t=-500 to 0 ms for response aligned data, see previous section 6.3.2 above) revealed significantly higher phase coherence in the trials that followed an error (cue-aligned: post error coherence=14.3±8.3%, post correct coherence =-5.2±5.2%, $p<0.05$, permutation testing; response aligned: post error coherence=13.3±6.8%, post correct coherence =-9.9±4.4%, $p<0.05$, permutation testing). As a control, we also analysed the theta band phase coherence between the STN LFP and the EEG over the parietal cortex (Pz-Cz) and found no error related differences (data not shown).

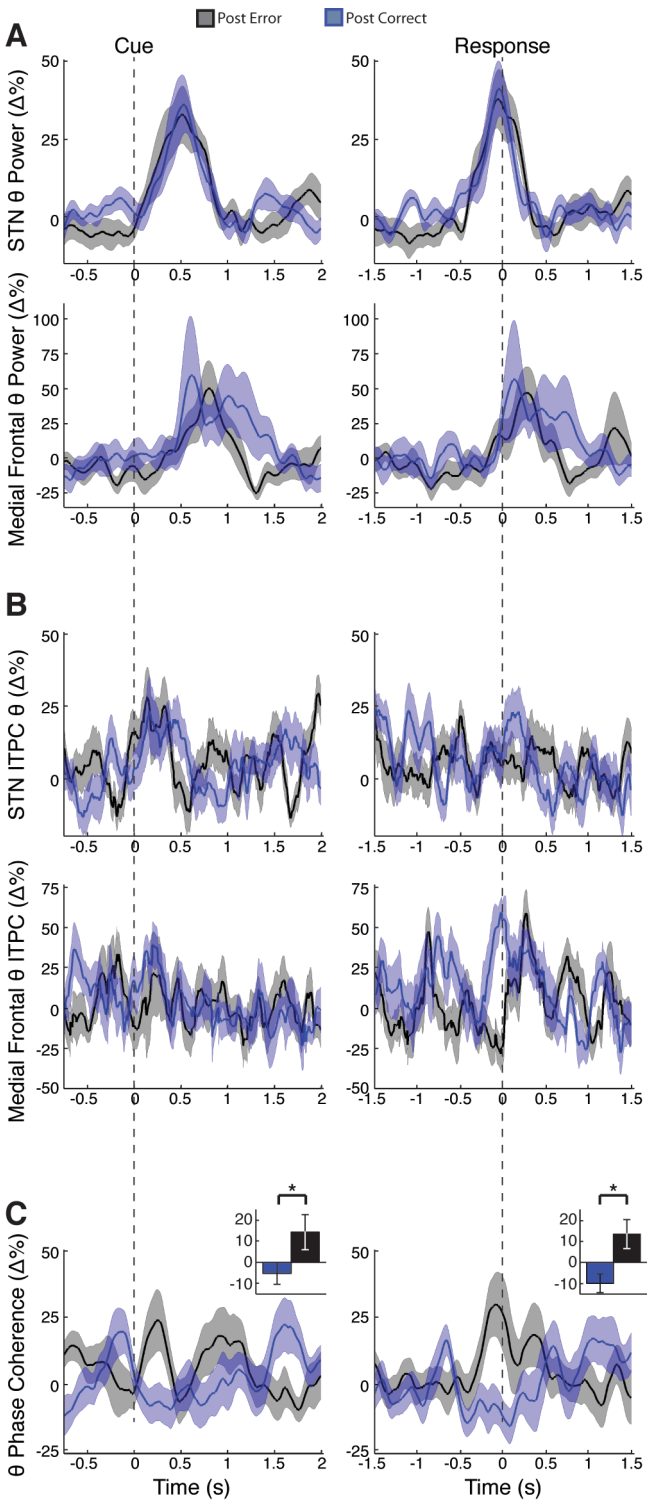


Figure 6.3. Theta STN-mPFC coherence is higher on trials that follow an error.
Same as **Figure 6.2**, but for the correct trials (black trace) that took place when an error was committed on the previous trial (not shown). Correct trials that followed a correct trial that was reaction time matched with an error trial are shown in blue. Insets in panel C show the group average ± SEM phase coherence in the 300ms following the cue (left) or the 500ms preceding the response (right). * denotes $p<0.05$.

6.3.4 Pre-stimulus correlates of conflict behaviour

As the previous section highlighted that task related activity on one trial can be related to electrophysiological activity on the subsequent trial, we decided to further explore the post-response, pre-subsequent warning cue STN power differences reported in Figure 6.2A. When we analysed the pre-stimulus periods (t=-2000 to 0 ms) that followed high and low conflict trials, we observed a significant difference in power approximately midway through this period (Figure 6.4A). This difference did not reflect a "spill over" of the conflict related differences of the previous trial as there was a period in between the pre-response differences of the previous trial and the pre-stimulus differences on the subsequent trial in which the power had returned to baseline levels for both conditions. Moreover, when we median split high conflict trials into two populations based on whether they were in the fastest half or the slowest half of the high conflict trials, we observed significant differences in the pre-warning cue power levels between the fastest and slowest high conflict trials (Figure 6.4B, bottom). During the 1.25 s that preceded the warning cue, the power level for the slowest high conflict trials was 2.8±0.7% higher than the average baseline power observed during all trials, and the power level for the fastest high conflict trials was 1.7±0.6% lower than the average baseline ($p<0.001$, permutation testing). This effect was not present during the low conflict trials (Figure 6.4B, top), suggesting that the pre-stimulus differences only affect reaction time when the subsequent stimulus contains conflict. Further support for this claim stems from our finding that the pre-stimulus power levels significantly correlated with trial reaction time in the high conflict condition (mean $R=0.10\pm0.03$, $p<0.05$, one-sample t-test), but not in the low conflict condition (mean $R=-0.01\pm0.03$, $p=0.6$, one-sample t-test; high vs. low $p<0.05$, permutation testing Figure 6.4C).

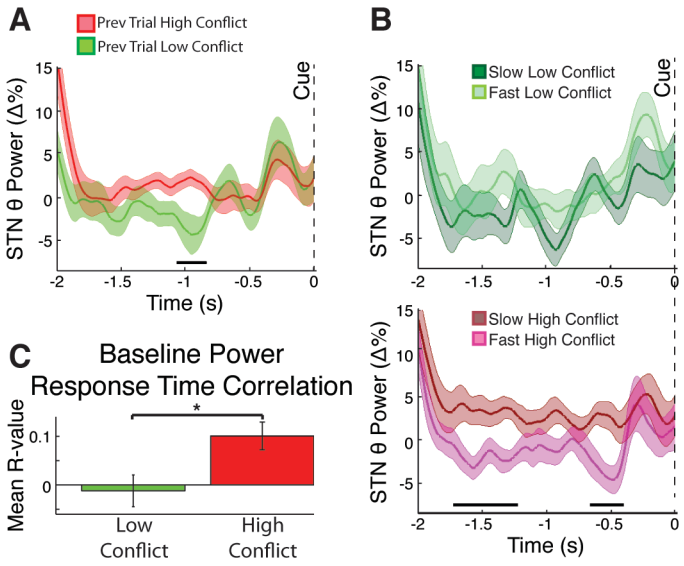


Figure 6.4. Inter-trial STN theta power is related to slower reaction times during conflict.
(A) The STN theta power (group average ± SEM, n=21) during the time periods that followed low and high conflict trials is shown aligned to the onset of the arrows for the subsequent trial (t=0). High conflict trials were followed by significantly higher theta power during the "baseline" of the next trial. Black line denotes significant time points ($p<0.05$ after correcting for multiple comparisons, permutation test). **(B)** A median split of the low and high conflict trials based on reaction time revealed that the slowest high conflict trials had higher pre-warning cue theta power than the fastest high conflict trials. **(C)** High conflict trial reaction time significantly correlated with the theta power that preceded the warning cue onset (t=-1.75 to -.5, relative to arrow onset at t=0). This effect was not present during low conflict trials. Across-subject average of within-subject Spearman correlation coefficients is shown for both conditions (average ± SEM). * denotes $p<0.05$.

6.4 Discussion

The results we present here corroborate previous findings we and others have published concerning the potentially crucial roles of theta oscillations and of the STN during conflict (Cavanagh et al., 2011; Zavala et al., 2013, 2014). Our results also provide several novel insights concerning the relationships between the STN, mesial frontal cortex, and cognitive control. High conflict trials demonstrated higher theta phase coupling between the mesial frontal EEG and the STN, and these differences occurred early in the trial before any power differences took place. Furthermore, error trials were followed by increased phase coherence between the mesial frontal EEG and the STN on the subsequent trial. Finally, pre-stimulus differences in STN theta power correlated with the reaction time of the subsequent trial, but only when that trial contained conflict. Taken together, these data help explain why disruption of this network by DBS can influence impulsivity during conflict and suggest the hitherto untested hypothesis that DBS may

also interfere with the behavioural adjustments that should take place when mistakes are committed.

6.4.1 Conflict related STN theta power and phase play separate roles within a trial

The importance of mPFC-STN theta connectivity was first suggested by a study showing that DBS to the STN disrupts the relationships between mPFC theta power and conflict related changes in reaction time (Cavanagh et al., 2011). To our knowledge, however, only the work described in Chapter 4 has directly shown increased coherence between the mPFC and the STN during conflict (Zavala et al., 2014). The latter study used a gradually adapting dot motion task to show that trials containing slow increases in conflicting information were associated with increased mPFC-STN coherence and of mPFC drive of STN theta activity. Here, we have reproduced the increase in mPFC-STN phase coherence using the flanker task, which involves both rapid onset conflict and ITPC increases. In contrast to the phase coherence, though, we did not observe any conflict related differences in ITPC during the task. This suggests that only when phases realign in a way that produces a specific and sustained phase difference between the two structures does activity in the two structures become coherent. The finding may also help explain previous contradictory results between studies showing task related differences in ITPC in the mPFC (Cavanagh et al., 2012) and STN (Zavala et al., 2013, Chapter 3 and Chapter 5 above), and others that do not (Cohen and Cavanagh, 2011; Zavala et al., 2014). It is the relative phase difference between mPFC and STN activities that is important, not the absolute phase in each.

The rapid conflict onset paradigm we used also allows us to make claims regarding the relative timing of theta activity. One unexpected finding in our data was that the STN LFP theta power increase actually occurred before that of the mesial frontal EEG. Indeed, the mesial frontal theta increase peaked after the response. Given that we have previously shown that mesial frontal oscillations drive those in the STN during conflict, this finding would seem paradoxical. A potential resolution to this paradox is provided by the phase of the theta oscillations. Unlike the

theta power increases, both the mesial frontal EEG and the STN LFP ITPC increases occurred at the same time, early during the trial, and it was these time periods that were associated with increased mesial frontal-STN coherence. The late increases in mesial frontal and STN theta power may therefore reflect other activity besides whether or not conflict is present. Indeed, a recent study by Cohen and Van Gaal (2014) showed that increases in mesial frontal theta power are associated with EMG detected "partial" errors in which the incorrect response is partially activated in a subthreshold manner. Interestingly, theta power differences between correct trials and partial error trials did not take place until after the partial error began. In light of this and our findings, we propose that early changes in mesial frontal-STN theta band coupling may be responsible for delaying all responses only during conflict, while late changes in theta band activity may be responsible for suppressing only the incorrect response during all trials. This hypothesis may help explain why low conflict trials also show a late increase in STN theta power as well as a late increase in mesial frontal-STN phase coherence. According to this interpretation, the higher increases in theta power that occur late during high conflict trials may reflect a greater drive needed to inhibit the incorrect response due to a stronger activation of that response by the flanking arrows.

6.4.2 STN theta activity also plays a role across trials

Another key finding from this study is that mesial frontal-STN connectivity seems to be involved in post error monitoring, which is thought to be one of the core functions of the mPFC (Ridderinkhof et al., 2004; Danielmeier and Ullsperger, 2011). It has previously been shown that error trials are associated with increased mPFC BOLD activity (Carter et al., 1998; Kerns et al., 2004) and theta power (Cavanagh et al., 2009, 2012). Furthermore, mPFC theta power is thought to underlie post error slowing (Rabbitt, 1966)(Rabbitt, 1966), either by interactions with DLPFC (Kerns et al., 2004; Hanslmayr et al., 2007; Cavanagh et al., 2009; Cohen and Cavanagh, 2011; Cohen and van Gaal, 2014) or by adjusting the excitability of the motor cortex following errors

(Danielmeier et al., 2011). Though it has been shown that mPFC error related theta activity can alter firing activity in the motor cortex (Narayanan et al., 2013), the mechanism by which it could communicate with the motor cortex was previously unknown as the two sites are not robustly connected (Wise, 2008). Here, we show that the STN provides a path by which mPFC theta oscillations might influence the excitability of the motor cortex following errors and thus potentially allow for error-related behavioural adjustments. Nevertheless, in our paradigm there was no slowing of reaction time in trials that followed an error, perhaps because of our use of an inter trial interval that is too long for post error slowing in healthy subjects (Jentzsch and Dudschig, 2009; Danielmeier and Ullsperger, 2011) or possibly because mesial frontal-STN connectivity might be impaired in our patient group, either as a function of PD or due to temporary stun effects at the level of the STN following surgery (Mann et al., 2009).

Finally, we observed that high conflict trials were followed by elevated STN LFP theta power levels in the baseline period before the subsequent trial, and that the reaction time in high conflict trials positively (albeit, weakly) correlated with that trial's baseline theta power. Interestingly, the baseline theta levels did not seem to influence low conflict trial reaction time, which is in line with other studies showing that some theta activity only correlates with behaviour during conflict (Cavanagh et al., 2011; Oehrn et al., 2014). Though our results are consistent with the posited braking effect of conflict related STN theta activity, they are in disagreement with what would be expected in the context of the Gratton effect (Gratton et al., 1992). That said our paradigm elicited no Gratton effect on reaction time here or previously (Zavala et al., 2013, Chapter 3 and Chapter 5), which may be related to impaired congruency sequence effects in PD (Rustamov et al., 2013).

Here we have investigated the mechanisms by which communication between brain regions may rapidly influence behaviour. However, the picture presented is likely to be incomplete, as many studies have shown that mPFC-IFC interactions are also related to within trial conflict processing

and across trial adjustments to conflict and errors (Kerns et al., 2004; Hanslmayr et al., 2007; Cavanagh et al., 2009; van Driel et al., 2012; Cohen and van Gaal, 2014; Oehrn et al., 2014). Nevertheless, our study suggests that theta activity and its synchronisation between cortical and subcortical structures may play a role in conflict and error related adaptations. Though disruption of this theta activity by either DBS (Cavanagh et al., 2011) or dopaminergic medication (Rodriguez-Oroz et al., 2011) has been shown to influence impulsivity, it remains to be determined if DBS affects across trial adjustments in behaviour or whether other disorders involving impaired decision-making and poor impulse control demonstrate altered theta activity in these networks (Fitzgerald et al., 2005; van Meel et al., 2007).

Chapter 7. Discussion

The studies reported in this thesis provide several lines of evidence that the STN is indeed involved in response inhibition during decisions involving action selection in the face of conflict. STN theta power increases during action selection, and these increases are higher during conflict (Chapters 3-6). The theta phases of STN oscillations re-align themselves across trials during action selection (chapters 3, 5, 6), possibly to prepare the motor system for the possibility of conflict. During conflict, the STN theta phases become coherent with those of the medial pre-frontal cortex (mPFC, chapter 4, 6), a structure that has often been implicated in cortical monitoring during conflict. This increased synchrony between the STN and the mPFC may underlie the entrainment of individual STN neurons to the theta rhythm specifically during conflict (chapter 5). The theta-spike phase locking was associated with increased response times, with altered firing rates, and with additional spike-phase entrainment in the beta frequency band, which has often been implicated in anti-kinetic inhibition. Finally, the STN may also be involved in other cognitive processes such as across trial adaptations to conflict and errors. The pre-stimulus period of a trial that followed conflict demonstrated increased theta power, and the pre-response period of a trial that followed an error demonstrated increased theta phase coherence between the mPFC and the STN.

7.1 Potential limitations of this thesis

Before discussing how this work fits in with previous hypotheses of STN function, it is important to discuss several limitations of this thesis. First and foremost is the fact that all of the recordings shown here were obtained, out of necessity, from the STN of patients with severe Parkinson's disease. Thus the generalizability of these results to the healthy human basal ganglia must be interpreted with caution (Williams et al., 2002; Hammond et al., 2007). That being said, I believe there are several lines of evidence suggesting that the activity we observed is indeed very similar to what would be observed in a healthy population. First, despite the impaired motor control that

was consistently present in our population of subjects, these subjects showed similar reaction time delays during conflict as the ones that have been reported in healthy human subjects (Eriksen and Eriksen, 1974). Given that the behavioural effects of conflict were similar between PD and healthy subjects, it would stand to reason that the electrophysiological effects would also be related. Second, many of the electrophysiological results that have been studied in the basal ganglia have subsequently been replicated in healthy non-human primates or in other DBS patient populations that do not have dopaminergic deficits such as those suffering from obsessive-compulsive disorders (OCD). For example, the presumably anti-kinetic rapid increases in STN beta band activity that were found while PD patients performed the Go-NoGo and stop signal reaction time task (Kühn et al., 2004; Ray et al., 2012; Alegre et al., 2013) have been subsequently reproduced in primates (Leventhal et al., 2012) and in patients with OCD (Bastin et al., 2014). Given that most of the literature concerning the role of mPFC theta oscillations during conflict has focused on healthy humans (Cavanagh et al., 2011, 2012; Cohen and Cavanagh, 2011), the increased coherence we observed between the STN and the mPFC fits perfectly well with what would be expected given their anatomical connectivity (Smith et al., 1998; Nambu et al., 2002; Haynes and Haber, 2013), imaging work related to their coactivation during response inhibition (Aron et al., 2007; Jahfari et al., 2012; Herz et al., 2014), and computational modelling work (Wiecki and Frank, 2013). Third, in all of the studies besides those recorded intra-operatively (Chapter 5), the patients performed the task while on their regular anti-parkinsonian medication in an attempt to study normal basal ganglia activity to the greatest extent possible. Given that the intra-operative study showed similar theta power increases as the post-operative studies conducted on medication, I suspect that the conflict related theta spike phase locking and beta band entrainment were not related to the lack of dopamine (although see section 7.5 "Paradoxes of deep brain stimulation of the STN" below for a discussion on how dopamine may alter theta activity in specific PD populations with L-Dopa induced impulsivity).

A second limitation of this study that should be kept in mind is that LFP signals can be subject to volume conduction. Thus the effects we report could reflect activity from structures near to, but outside of the STN. To mitigate this effect, all electrophysiological analyses were conducted on re-referenced electrodes i.e. by subtracting from each electrode the activity of an adjacent electrode (Chapter 3, Chapter 4 and Chapter 6) or the activity of all electrodes inside the basal ganglia (Chapter 5). Though this would greatly decrease the impact of volume conduction on our results, one drawback of this method is that it diminishes the spatial resolution of our recordings. This may have contributed to the third limitation of our study: we never localized any of our effects to specific subdivisions of the STN. Several previous researchers have argued that the STN may have spatially segregated areas dedicated to motor processing (dorsal areas) and to more cognitive functions in the associative and limbic domains (ventral areas) (Rodriguez-Oroz et al., 2001, 2011; Mallet et al., 2007; Lardeux et al., 2009). Other studies, however, have suggested that the STN is not neatly segregated into different domains, and neurons throughout the nucleus can be involved in motor control and cognition (Haynes and Haber, 2013; Eisenstein et al., 2014). As we did not want to use any specific channel selection criteria that could potentially bias our results (i.e. many studies use the channel with the highest beta as the channel that is "most" representative of STN activity), we chose to include all STN recordings in our analysis and average across all the contacts within one STN. Nevertheless, when we did attempt to compare the activity of contact pairs with the highest beta band (used as a metric to find the electrode in the motor, dorsal area of STN) to the remaining activity, we observed no differences in theta band reactivity (chapters 3-4). This, however, does not necessarily suggest that the STN is homogenous. Rather, it implies that without high resolution imaging, it is difficult to infer the precise location within the STN using only simple biomarkers such as theta and beta power.

The final limitation of our studies that should be kept in mind is that most of the work I have presented is correlative in nature (the one exception being the on vs. off DBS stimulation during

the flanker task, which resulted in higher error rates) and must be corroborated with causative interventions that are now possible through transcranial alternating current stimulation (TACS, Pogosyan et al., 2009; Joundi et al., 2012) and frequency specific DBS (Chen et al., 2011; Cagnan et al., 2013). Furthermore, the low signal to noise ratios of LFP and EEG data required that I use very crude methods to test differences in electrophysiological activity. This involved averaging across all trials within a condition and comparing this to the average of all trials in a different condition. Other work has suggested that not all "high conflict" trials necessarily induce conflict in the brain, and not all "low conflict" trials are necessarily free from the conflict that comes with doing a task in which some trials do involve conflict (Yeung et al., 2011). Our best attempt to separate the "high conflict" trials based on amount of conflict experienced in a given trial was to median split trial groups based on the reaction time. We also attempted on a few occasions to correlate individual trial activity to behaviour. Though these correlations did reveal significant associations between electrophysiology and behaviour at the single trial level, most of the correlations we observed were quite week (around the order of r=0.1). A large part of this is likely related to the less than ideal signal to noise ratios as well as the attenuation of local power reactivity following DBS surgery (i.e. the "stun effect", Chen et al., 2006). The low correlations, however, suggest that the STN is only one of several structures the brain uses to influence reaction time. Nevertheless, the techniques we employed revealed consistent, significant differences between the activity of the STN during low and high conflict trials, which allows me to make some inferences about how the STN may be involved in cognition.

7.2 A hypothesized role of STN theta band power during action selection

In the discussion of Chapter 3 I highlighted several possible roles of STN theta band activity that would be consistent with the higher theta power levels observed during conflict. One possibility I proposed was that theta activity could be pro-kinetic in nature and reflect the sum of all evidence

supporting the selection of several actions. According to this interpretation, the higher theta power observed during high conflict trials could reflect the higher amounts of "evidence" inputs sent to the STN during conflict due to the simultaneous activation of the correct and the incorrect responses. Our subsequent work (Chapter 4) with the gradually adapting congruent dots experiment, however, strongly suggests that this interpretation is incorrect. In the dots task, the conflict was present early in the trial, but gradually disappeared with time. During the high conflict periods of this task, the theta power was elevated relative to the power observed during the same time points during low conflict trials. It was only once the theta power decreased back to baseline levels that a response was executed. This activity is more in line with an anti-kinetic interpretation of theta band activity as previously proposed by others (Cavanagh et al., 2011). An anti-kinetic interpretation of theta power, however, seems to be at odds with the fact that STN theta power during the low and the high conflict trials of the flanker task continued to rise until a response was executed. Thus, our results in the flanker task beg the question, what is the STN theta power doing? We propose that the theta power increases observed in the STN are anti-kinetic in nature and are responsible for inhibiting any incorrect actions. During conflict, the higher theta power at the time of a response reflects a greater need to suppress an incorrect action (relative to a low conflict trial) because that action has received additional activation by the flanking arrows. This hypothesis is in line with the diffuse projections observed in the STN (Hazrati and Parent, 1992), with the anti-kinetic role posited for cortical theta activity (Cavanagh et al., 2011, 2012), and with the higher STN theta power observed during conflict (Chapters 3-6 and Cavanagh et al., 2011; Fumagalli et al., 2011; Brittain et al., 2012; Zavala et al., 2013, 2014). Finally, an anti-kinetic interpretation of theta band activity is also consistent with the higher STN theta power levels we observed during the pre-stimulus period of the slowest high conflict flanker trials.

7.3 A hypothesized role of STN theta phase during action selection, conflict, and post error monitoring

In contrast to the theta power increases that peaked at the time of the response, all of the changes in theta band ITPC we observed were strongest at the onset of an event such as the presentation of the flanker stimuli. As the theta phase resets occurred during low and high conflict trials as well as during the presentation of the warning cue (Chapter 3, Chapter 5 and Chapter 6), we propose that they may be related, at least in part, to an orientation response preparing subjects for the actions that may be required soon after the stimulus. This claim is bolstered by the finding that the patients with the weakest ITPC response to the stimulus (calculated across all trials) had the slowest reaction times (Chapter 3).

The fact that the ITPC increases occurred at similar time points in the mPFC and in the STN suggests that they may also be involved in preparing different cortical and subcortical structures to become coherent if need be. In line with this proposal, the phase coherence results (Chapter 4 and Chapter 6) suggest that indeed the high conflict trials of the flanker task and the dots task were associated with increases in synchrony between the mPFC and the STN. This finding may help explain why we occasionally found contradictory effects when we compared ITPC values for low and high conflict trials. In chapter 3 we presented no difference in ITPC between low and high conflict trials, but we did see that the slowest half of the high conflict trials had lower ITPC values. In contrast, in chapter 5 we found higher ITPC values during the high conflict trials, and in chapter 6 we observed no difference between low and high conflict trial ITPC values. Perhaps it is the relative phase between the mPFC and the STN that is crucial for response inhibition, not the absolute phase of the STN alone. The increased synchrony we observed is in line with the anti-kinetic interpretation of theta band activity as well as with the presumed role of cortical monitoring assigned to conflict related mPFC activity (Cavanagh et al., 2012). In light of the theta band spike phase locking results I presented in chapter 5, I propose that when the mPFC and the

STN are synchronized, the phase of the theta band is able to entrain STN neurons that are activated early during the task. Perhaps it is this entrainment to theta band activity that results in the altered firing rates observed in the STN cells during high conflict (chapter 5) which then inhibit inappropriate responses long enough for other structures to select the correct action.

7.4 The STN is a point of convergence in a larger response inhibition network

The proposal that the mPFC could detect conflict and activate STN neurons is consistent with previous studies. Many imaging and event related potential EEG studies demonstrate higher mPFC activity during conflict (Botvinick et al., 2004). Moreover, neurons in two key structures within the mPFC, the ACC and the pre-SMA, show STN-like changes in discharge rate during response inhibition and stopping. Isoda and Hikosaka extended their saccade priming task described in section 1.4.1 "Firing rate studies of single neurons" to record the firing rates of non-human primate neurons in the STN and pre-SMA during trials that involved conflict between a pre-potent and a controlled saccade (Isoda and Hikosaka, 2007, 2008). Both structures showed similar firing rate increases during correct trials in which the pre-potent response was successfully inhibited. More recently, others have shown that individual neurons of the human ACC discharge at a higher rate during conflict, and that this correlates with the amount of conflict as it does in the STN (Sheth et al., 2012; Zaghloul et al., 2012).

In light of the evidence highlighting the role of the mPFC during conflict, Frank has recently revised his "hold your horses" model of the basal ganglia's role in decision-making to include various cortical structures as well (Wiecki and Frank, 2013). In this model, the ACC detects conflict and subsequently activates the STN to inhibit responses. The work I have presented here lends strong support to this model and suggests that the theta band may play a crucial role in connecting these two structures during conflict. Our work, however, also leaves several questions

unanswered, such as what role the conflict related beta band spike-phase entrainment may play (Chapter 5) or how the mPFC and the STN selectively synchronize their theta band activity only during conflict given that both structures show ITPC increases when low and high conflict flanker trials begin. Wieke and Frank's model may help shed light on these outstanding questions. In their model, they posit that the inferior frontal cortex (IFC) also plays a role in response inhibition, and all three structures (the STN, the mPFC, and the IFC) work together to alter the dynamics of action selection. Support for this model stems from the fact that all three nodes of the stopping network are connected and are activated in a similar manner during response inhibition (Aron et al., 2007).

As of yet, however, only a handful of studies have analysed the mechanisms by which the mesial and lateral frontal cortices interact with each other and with the STN during conflict. In addition to our studies showing increased mPFC-STN synchrony during conflict (Zavala et al., 2014, Chapter 4 and Chapter 6), other studies have shown frequency specific synchrony between the two cortical nodes during cognitive processes. High conflict and error trials in the flanker and Stroop tasks are associated with increased theta band synchrony between the mPFC and the IFC (Hanslmayr et al., 2007; Cavanagh et al., 2009; Cohen and Cavanagh, 2011; Oehrn et al., 2014). Beta oscillations also seem to link the mPFC and the IFC during the stop signal reaction time task (Swann et al., 2012), and beta band synchrony between these cortical structures increases just before errors are committed by musicians playing the piano (Ruiz et al., 2011). The authors of the latter study suggested that the beta synchrony between the two cortical structures reflected stopping activity trying to prevent errors from occurring, and this was supported by the fact that errors were played at lower volumes. Though there has yet to be a study showing increased coherence between the STN and the IFC, it is possible that beta oscillations play a role in synchronizing the two structures. During the stop signal reaction time task, similar beta increases

as those seen in the STN during this task (Ray et al., 2012) have also been observed over the IFC (Swann et al., 2009, 2011, 2012).

The joint role that theta and beta oscillations play in connecting the mPFC and the IFC during response inhibition, together with the similar beta reactivity profiles of the STN and IFC during abrupt stopping may help explain the beta band spike phase locking we observed that was specific to high conflict trials (Chapter 5). According to the Wieke-Frank model of cortical-subcortical interactions during conflict and response inhibition, the ACC may be more heavily involved in conflict related response delays, while the IFC is more involved in global inhibition of all responses. Therefore one explanation of our data is that theta oscillations link the mPFC and the STN directly during conflict, but when the inhibition must be stronger, IFC inputs in the beta band provide additional anti-kinetic activity to the basal ganglia via the STN. Thus even though most of the structures in the basal ganglia show decreases in beta activity during action selection (Leventhal et al., 2012), rapid hyperdirect inputs of beta activity from cortical structures to the STN may further delay movements during conflict. The increased synchrony that others have observed between the mPFC and the IFC during response inhibition may allow the two cortical nodes to recruit each other to further activate the STN when one of the cortical nodes alone is unable to fully inhibit incorrect actions. Indeed one recent study (Herz et al., 2014) used transcranial magnetic stimulation (TMS) to show that when the activity of the mPFC is impaired by repetitive TMS, the fMRI determined effective connectivity between the STN and the IFC increased to enhance response inhibition. Though it remains to be determined why silencing mPFC would lead to enhanced inhibition, the authors suggest that the strengthening of IFC-STN connectivity was compensatory in this paradigm. In light of their findings, the authors endorsed the Wieke-Frank model of response inhibition, where the mPFC and the IFC activate the STN in parallel, not in series. According to this interpretation, the theta and beta bands could provide

frequency specific channels of communication between the cortical structures and the STN (Fries, 2005).

Regardless of the exact organization underlying the response-inhibition network, the expanding literature on the role of oscillatory synchronization within the network suggests that the traditional firing rate models of basal ganglia connectivity are incomplete (see section 1.1.4 "Beyond the rate model"), as they do not include these important phenomena. This offers a potential explanation for a neglected shortcoming of the original 'hold-your-horses' model. As discussed earlier, the core evidence supporting this model is that high frequency DBS of the STN induces a degree of impulsivity, interpreted as resulting from a blocking of conflict related increased firing rates in the STN. Yet, the prevailing evidence now suggests that high frequency DBS actually drives STN output (Windels et al., 2000; Hashimoto et al., 2003; Stefani et al., 2005; McIntyre and Hahn, 2010), which should theoretically increase response caution. However, DBS also seems to block or replace patterned output (Eusebio et al., 2011; Moran et al., 2012; Whitmer et al., 2012) so the effects of high frequency DBS on key changes in physiological oscillatory synchronization during response inhibition, rather than firing rates, may reconcile the hold-your-horses model with emerging views of DBS effects.

7.5 Paradoxes of deep brain stimulation of the STN

Our results suggest that the STN plays at least some part in inhibiting impulsivity by delaying the selection of in appropriate response. Indeed there are cross sectional studies demonstrating that DBS to the STN makes patients more impulsive (Hälbig et al., 2009) in their day to day lives. Most individuals who receive DBS, however, do not suffer from pathological impulsivity following electrode implantation. This would argue against the importance of the STN in response inhibition and suggests that there are also other structures in the brain that are able to inhibit impulsivity. It is important to remember, however, in what context we have argued that the

STN is important. Almost all of the tasks we and others have used to study the STN's role in conflict involve very rapid decisions that require quick response inhibition to prevent impulsive responding or to overcome natural biases that would favour the incorrect response. This is why most of the error trials that occur in these tasks had much faster reaction times than the correct trials; the inhibition most likely did not arrive soon enough to prevent the execution of the pre-potent response (Brittain et al., 2012; Schmidt et al., 2013). We propose that those studies suggesting that STN DBS impairs response inhibition can be reconciled with the clinical reality that pathological impulsivity is only very infrequently seen as a complication of STN DBS, if the response inhibition exerted through the STN is only rapid and not sustained. Indeed, there is evidence that STN DBS only disrupts rapid inhibition that occurs early on in a trial (Wylie et al., 2010). By examining the behaviour of DBS subjects performing the Simon task (Simon, 1969), Wylie et al were able to show that DBS induces opposing effects depending on the type of inhibition being activated (Wylie et al., 2010). During trials with very fast response times, they observed decreases in accuracy consistent with the previously described finding that disrupting the STN leads to quick, impulsive errors (see section 1.3 "Lesion and DBS studies show impaired response inhibition during conflict" above). During the slowest trials, however, they observed that DBS actually leads to improvements in response inhibition by suppressing the effect of the incorrect response. Thus, STN related inhibition is important in paradigms like Go-NoGo, Stop, Stroop, flanker, and even the probabilistic reward task where a response is required in each trial. However, many important decisions, whether made over conflicting choices or not, do not require an immediate action. Such delayed actions would not be affected by the type of rapid response inhibition associated with the STN and yet may be it is precisely these delayed action decisions that are affected by the kind of pathological impulsivity that negatively impacts work and relationships. The probable lack of STN involvement in these decisions possibly explains why most DBS patients are able to continue their day-to-day lives without any noticeable signs of pathological impulsivity. Nevertheless, the higher impulsivity scores reported by DBS patients in

a questionnaire probing day-to-day impulsivity, albeit in a cross-sectional study (Hälbig et al., 2009), and the rare cases of DBS triggered pathological impulsivity (Romito et al., 2002; Herzog et al., 2003; Lu C et al., 2006; Mandat et al., 2006; Smeding et al., 2007) still implicate the STN in the regulation of impulsive actions.

Another outstanding question relates to the fact that not all experimental studies demonstrate that STN DBS impairs response inhibition, and why some even demonstrate DBS related improvements in inhibition (Schroeder et al., 2002; van den Wildenberg et al., 2006; Mirabella et al., 2012). Factors such as task details may again be important, as may be the precise targeting given that the effects of DBS likely vary according to the specific subsections of the STN that are stimulated (Mandat et al., 2006; Mallet et al., 2007; Rodriguez-Oroz et al., 2011; Alegre et al., 2013). Additionally, patient state also seems to play an important role in determining the outcome of DBS on response inhibition. Ray et al observed that DBS had opposing effects on the stop signal reaction time task when the population of DBS patients was divided into two sub populations (Ray et al., 2009). The patients that had a similar baseline performance as healthy control subjects showed a detrimental effect on stopping when their stimulators were turned on. Patients who were already bad at the task at baseline, however, showed an improvement in response inhibition with DBS. One possible explanation for this may be related to DBS induced modulations of other nodes in the network (Schroeder et al., 2002; Swann et al., 2011). This may possibly explain why three other studies have also reported that stop signal reaction time improves with DBS (van den Wildenberg et al., 2006; Swann et al., 2011; Mirabella et al., 2012). One of these studies (Swann et al., 2011) related the improvement to enhanced theta and beta power increases over frontal cortex (particularly that on the right-side) in successful stop trials made during DBS, leading to the testable hypothesis that DBS patients who experience improvements in response inhibition may show compensatory increased coupling within the response inhibition network.

A more in depth knowledge of the STN and its connections may also help shed light on recent unexpected findings related to the electrophysiology of patients with impulse control disorders that predate surgery. Rosa et al., recently showed that during a gambling task, Parkinsonian patients with pathological gambling adopted a riskier strategy than patients without pathological gambling (Rosa et al., 2013). Those that adopted a conservative strategy showed no difference in low frequency STN power (2-12 Hz) between low and high conflict trials. Those that adopted the risky strategy, on the other hand, showed differences between the trial types due to lower amplitude low frequency power during the low conflict trials. Furthermore, in a study comparing the baseline STN activity of Parkinsonian patients with impulse control disorders to the activity of Parkinsonian patients without such symptoms, Rodriguez-Oroz et al. found that the impulsive patients had levodopa-induced increases in theta oscillations in ventral STN as well as in fronto-STN theta coherence (Rodriguez-Oroz et al., 2011). These studies seem to call into question the idea that theta might delay or inhibit responses during conflict. Viewing theta as an anti-kinetic conflict related signal as encapsulated in current computational models would predict that impulsive patients should have decreased baseline theta levels. One way these findings could be reconciled is if the abnormal theta activity observed in the STN of impulsive patients was not responsible for the symptoms, but rather reflected a breakdown of the relationship between theta and behaviour or even a compensatory mechanism that ameliorates the deficits induced by chronic dopaminergic activation. The latter possibility may also help to explain why STN DBS per se, does not seem to treat impulse control disorders, as would be expected if the exaggerated STN theta activity were primarily pathological, and can even, on rare occasions, induce them. Rather it is the associated decrease in dopaminergic medication with chronic STN DBS that leads to improvement in such impulse control disorders (Castrioto et al., 2014). Together these observations prompt the radical view that interventional therapies should seek to promote local theta and perhaps even beta (Joundi et al., 2012) in the STN so as to ameliorate conditions characterized by inadequate response inhibition.

7.6 Future Perspectives

The work I have presented here indicates that the STN is an important player in an extensive response inhibition network. Our correlative evidence suggests that the STN's operation within this distributed network may involve the graded, task and context-dependent manipulation of oscillatory activities, particularly in the theta and beta frequency bands. However, as most of the results we have presented only reveal a correlative relationship between electrophysiology and response inhibition, it remains for existing evidence to be complemented by interventional studies demonstrating causal links between oscillatory activities and response inhibition in the frontal cortical-STN network. Though I have already attempted to use Transcranial Alternating Current Stimulation (TACS) to induce theta oscillations over the mPFC during the flanker task, we were unable to induce conflict related slowing (data not shown). In the future I may repeat these experiments using high resolution TACS (Minhas et al., 2010) that will allow us more precise control of the targets that are stimulated and the currents that are delivered. Furthermore, our work has highlighted the need to understand the dynamics of the response inhibition network as a whole, as the STN cannot be considered in isolation. In the future, this will be facilitated by simultaneous electrical recordings from STN, mPFC, and IFC during decision-making and response inhibition tasks. We have already begun collecting intra-operative recordings from the macroelectrodes in the STN and from electrocorticography electrodes over the mPFC during a task that involves both high conflict trials and trials that require outright stopping. We also plan on recording from the mPFC and IFC of epilepsy patients while they perform the same task in an attempt to test the hypothesis that the theta and beta bands are involved in slightly different forms of response inhibition, both of which involve the STN. Finally, it remains to be determined how the activity of the STN propagates to its downstream targets. We have repeated the flanker task while recordings were acquired from the GPi, but we are still waiting for a few more patients to complete this data set. We have also begun recordings from the thalamus while patients perform a serial reaction time task that occasionally requires the inhibition of a learned sequence pattern.

As it stands, however, our existing data lead to a working hypothesis. Key has been the observation that the theta band synchronizes the activity of the STN with cortical structures implicated on executive function. This, coupled with the growing evidence that high frequency DBS may act by desynchronizing neural activity (Eusebio et al., 2011; Moran et al., 2012; Whitmer et al., 2012), would be compatible with DBS leading not only to a suppression of pathological oscillatory synchrony (Hammond et al., 2007), but also of physiological synchronization in the theta and beta bands during conflict. Accordingly, future research focused on preserving such physiological activity will likely prove to be instrumental in preventing cognitive side effects associated with DBS. Whereas DBS therapy for Parkinson's disease currently involves constant high frequency stimulation to the STN, future therapies are likely to involve more intelligent solutions that only stimulate when certain Parkinsonian biomarkers are present to a pathological degree. Already there have been two seminal studies showing that demand driven stimulation triggered by neuronal firing or excessively synchronized beta oscillations are at least as effective, if not better, than constant high frequency stimulation of the STN or its target the GPi (Rosin et al., 2011; Little et al., 2013). These intermittently active "closed-loop" studies therefore have the potential to reduce the disruptive effects of DBS on physiological oscillations that play a role in effective decision-making. Moreover, with the advent of DBS electrodes with greater numbers of contacts, steerable electrical fields and sensing capabilities, it may be possible to direct stimulation away from physiologically relevant signals like theta activity in the subthalamic region (Pollo et al., 2014). In the longer term, a richer understanding of the dynamics of the extended response inhibition network may also inspire novel, less invasive therapies. It has already been shown that non-invasive TACS over the motor cortex can affect the tremor of PD patients when it is delivered out-of-phase to the ongoing tremor activity (Brittain et al., 2013) and that beta band stimulation can greatly promote response inhibition in healthy subjects (Joundi et al., 2012). Future studies exploring the effects of TMS and TACS to the mPFC and IFC may reveal novel ways of influencing STN activity in PD, but

also in conditions where response inhibition is even more strikingly deficient, such as Gilles de la Tourette syndrome and obsessive compulsive disorders.

Chapter 8. References

Albin RL, Young AB, Penney JB (1989) The functional anatomy of basal ganglia disorders. Trends in Neurosciences 12:366–375.

Alegre M, Labarga A, Gurtubay I, Iriarte J, Malanda A, Artieda J (2003) Movement-related changes in cortical oscillatory activity in ballistic, sustained and negative movements. Exp Brain Res 148:17–25.

Alegre M, Lopez-Azcarate J, Obeso I, Wilkinson L, Rodriguez-Oroz MC, Valencia M, Garcia-Garcia D, Guridi J, Artieda J, Jahanshahi M, Obeso JA (2013) The subthalamic nucleus is involved in successful inhibition in the stop-signal task: A local field potential study in Parkinson's disease. Experimental Neurology 239:1–12.

Amirnovin R, Williams ZM, Cosgrove GR, Eskandar EN (2004) Visually Guided Movements Suppress Subthalamic Oscillations in Parkinson's Disease Patients. J Neurosci 24:11302–11306.

Antoniades CA, Bogacz R, Kennard C, FitzGerald JJ, Aziz T, Green AL (2014) Deep Brain Stimulation Abolishes Slowing of Reactions to Unlikely Stimuli. J Neurosci 34:10844–10852.

Aron AR, Behrens TE, Smith S, Frank MJ, Poldrack RA (2007) Triangulating a Cognitive Control Network Using Diffusion-Weighted Magnetic Resonance Imaging (MRI) and Functional MRI. J Neurosci 27:3743–3752.

Ballanger B, van Eimeren T, Moro E, Lozano AM, Hamani C, Boulinguez P, Pellecchia G, Houle S, Poon YY, Lang AE, Strafella AP (2009) Stimulation of the subthalamic nucleus and impulsivity: release your horses. Ann Neurol 66:817–824.

Bastin J, Polosan M, Benis D, Goetz L, Bhattacharjee M, Piallat B, Krainik A, Bougerol T, Chabardès S, David O (2014) Inhibitory control and error monitoring by human subthalamic neurons. Transl Psychiatry 4:e439.

Baunez C, Nieoullon A, Amalric M (1995) In a rat model of parkinsonism, lesions of the subthalamic nucleus reverse increases of reaction time but induce a dramatic premature responding deficit. J Neurosci 15:6531–6541.

Benabid AL, Pollak P, Hoffmann D, Gervason C, Hommel M, Perret JE, de Rougemont J, Gao DM (1991) Long-term suppression of tremor by chronic stimulation of the ventral intermediate thalamic nucleus. The Lancet 337:403–406.

Benabid AL, Pollak P, Louveau A, Henry S, de Rougemont J (1987) Combined (thalamotomy and stimulation) stereotactic surgery of the VIM thalamic nucleus for bilateral Parkinson disease. Appl Neurophysiol 50:344–346.

Benazzouz A, Gross C, Féger J, Boraud T, Bioulac B (1993) Reversal of Rigidity and Improvement in Motor Performance by Subthalamic High-frequency Stimulation in MPTP-treated Monkeys. European Journal of Neuroscience 5:382–389.

Benazzouz A, Piallat B, Pollak P, Benabid A-L (1995) Responses of substantia nigra pars reticulata and globus pallidus complex to high frequency stimulation of the subthalamic nucleus in rats: electrophysiological data. Neuroscience Letters 189:77–80.

Bergman H, Wichmann T, DeLong MR (1990) Reversal of experimental parkinsonism by lesions of the subthalamic nucleus. Science 249:1436–1438.

Bergman H, Wichmann T, Karmon B, DeLong MR (1994) The primate subthalamic nucleus. II. Neuronal activity in the MPTP model of parkinsonism. Journal of Neurophysiology 72:507–520.

Birkmayer W, Hornykiewicz O (1961) [The L-3,4-dioxyphenylalanine (DOPA)-effect in Parkinson-akinesia]. Wien Klin Wochenschr 73:787–788.

Bogacz R, Gurney K (2007) The Basal Ganglia and Cortex Implement Optimal Decision Making Between Alternative Actions. Neural Computation 19:442–477.

Botvinick MM, Cohen JD, Carter CS (2004) Conflict monitoring and anterior cingulate cortex: an update. Trends in Cognitive Sciences 8:539–546.

Brittain J-S, Probert-Smith P, Aziz TZ, Brown P (2013) Tremor suppression by rhythmic transcranial current stimulation. Curr Biol 23:436–440.

Brittain J-S, Watkins KE, Joundi RA, Ray NJ, Holland P, Green AL, Aziz TZ, Jenkinson N (2012) A Role for the Subthalamic Nucleus in Response Inhibition during Conflict. J Neurosci 32:13396–13401.

Brovelli A, Ding M, Ledberg A, Chen Y, Nakamura R, Bressler SL (2004) Beta oscillations in a large-scale sensorimotor cortical network: Directional influences revealed by Granger causality. PNAS 101:9849–9854.

Brown P (2003) Oscillatory nature of human basal ganglia activity: Relationship to the pathophysiology of Parkinson's disease. Mov Disord 18:357–363.

Brown P (2007) Abnormal oscillatory synchronisation in the motor system leads to impaired movement. Current Opinion in Neurobiology 17:656–664.

Brown P (2013) Making use of pathological synchrony in Parkinson's disease. Clinical Neurophysiology 124:834–835.

Brown P, Chen CC, Wang S, Kühn AA, Doyle L, Yarrow K, Nuttin B, Stein J, Aziz T (2006) Involvement of human basal ganglia in offline feedback control of voluntary movement. Curr Biol 16:2129–2134.

Brown P, Eusebio A (2008) Paradoxes of functional neurosurgery: Clues from basal ganglia recordings. Mov Disord 23:12–20.

Brown P, Mazzone P, Oliviero A, Altibrandi MG, Pilato F, Tonali PA, Di Lazzaro V (2004) Effects of stimulation of the subthalamic area on oscillatory pallidal activity in Parkinson's disease. Experimental Neurology 188:480–490.

Brown P, Oliviero A, Mazzone P, Insola A, Tonali P, Lazzaro VD (2001) Dopamine Dependency of Oscillations between Subthalamic Nucleus and Pallidum in Parkinson's Disease. J Neurosci 21:1033–1038.

Burbaud P et al. (2013) Neuronal activity correlated with checking behaviour in the subthalamic nucleus of patients with obsessive–compulsive disorder. Brain 136:304–317.

Burns RS, Chiueh CC, Markey SP, Ebert MH, Jacobowitz DM, Kopin IJ (1983) A primate model of parkinsonism: selective destruction of dopaminergic neurons in the pars compacta of the substantia nigra by N-methyl-4-phenyl-1,2,3,6-tetrahydropyridine. PNAS 80:4546–4550.

Buzsáki G, Anastassiou CA, Koch C (2012) The origin of extracellular fields and currents — EEG, ECoG, LFP and spikes. Nat Rev Neurosci 13:407–420.

Cagnan H, Brittain J-S, Little S, Foltynie T, Limousin P, Zrinzo L, Hariz M, Joint C, Fitzgerald J, Green AL, Aziz T, Brown P (2013) Phase dependent modulation of tremor amplitude in essential tremor through thalamic stimulation. Brain 136:3062–3075.

Carpenter RHS, Williams MLL (1995) Neural computation of log likelihood in control of saccadic eye movements. Nature 377:59–62.

Carp J, Kim K, Taylor SF, Fitzgerald KD, Weissman DH (2010) Conditional Differences in Mean Reaction Time Explain Effects of Response Congruency, but not Accuracy, on Posterior Medial Frontal Cortex Activity. Front Hum Neurosci 4:231.

Carter CS, Braver TS, Barch DM, Botvinick MM, Noll D, Cohen JD (1998) Anterior Cingulate Cortex, Error Detection, and the Online Monitoring of Performance. Science 280:747–749.

Cassidy M, Mazzone P, Oliviero A, Insola A, Tonali P, Lazzaro VD, Brown P (2002) Movement-related changes in synchronization in the human basal ganglia. Brain 125:1235–1246.

Castrioto A, Funkiewiez A, Debû B, Cools R, Lhommée E, Ardouin C, Fraix V, Chabardès S, Robbins TW, Pollak P, Krack P (2014) Iowa gambling task impairment in Parkinson's disease can be normalised by reduction of dopaminergic medication after subthalamic stimulation. J Neurol Neurosurg Psychiatry:jnnp – 2013–307146.

Cavanagh JF, Cohen MX, Allen JJ (2009) Prelude To And Resolution of an Error: EEG Phase Synchrony Reveals Cognitive Control Dynamics During Action Monitoring. J Neurosci 29:98–105.

Cavanagh JF, Sanguinetti JL, Allen JJB, Sherman SJ, Frank MJ (2014) The Subthalamic Nucleus Contributes to Post-error Slowing. J Cogn Neurosci 26:2637–2644.

Cavanagh JF, Wiecki TV, Cohen MX, Figueroa CM, Samanta J, Sherman SJ, Frank MJ (2011) Subthalamic nucleus stimulation reverses mediofrontal influence over decision threshold. Nat Neurosci 14:1462–1467.

Cavanagh JF, Zambrano-Vazquez L, Allen JJB (2012) Theta lingua franca: A common mid-frontal substrate for action monitoring processes. Psychophysiology 49:220–238.

Chen CC, Lin WY, Chan HL, Hsu YT, Tu PH, Lee ST, Chiou SM, Tsai CH, Lu CS, Brown P (2011) Stimulation of the subthalamic region at 20 Hz slows the development of grip force in Parkinson's disease. Exp Neurol 231:91–96.

Chen CC, Litvak V, Gilbertson T, Kühn A, Lu CS, Lee ST, Tsai CH, Tisch S, Limousin P, Hariz M, Brown P (2007) Excessive synchronization of basal ganglia neurons at 20 Hz slows movement in Parkinson's disease. Experimental Neurology 205:214–221.

Chen CC, Pogosyan A, Zrinzo LU, Tisch S, Limousin P, Ashkan K, Yousry T, Hariz MI, Brown P (2006) Intra-operative recordings of local field potentials can help localize the subthalamic nucleus in Parkinson's disease surgery. Experimental Neurology 198:214–221.

Cohen MX (2014) Fluctuations in Oscillation Frequency Control Spike Timing and Coordinate Neural Networks. J Neurosci 34:8988–8998.

Cohen MX, Cavanagh JF (2011) Single-trial regression elucidates the role of prefrontal theta oscillations in response conflict. Front Psychology 2:30.

Cohen MX, Gulbinaite R (2014) Five methodological challenges in cognitive electrophysiology. Neuroimage 85 Pt 2:702–710.

Cohen MX, Nigbur R (2013) Reply to "Higher response time increases theta energy, conflict increases response time." Clin Neurophysiol 124:1479–1481.

Cohen MX, van Gaal S (2014) Subthreshold muscle twitches dissociate oscillatory neural signatures of conflicts from errors. NeuroImage 86:503–513.

Coulthard EJ, Bogacz R, Javed S, Mooney LK, Murphy G, Keeley S, Whone AL (2012) Distinct roles of dopamine and subthalamic nucleus in learning and probabilistic decision making. Brain 135:3721–3734.

Courtemanche R, Fujii N, Graybiel AM (2003) Synchronous, Focally Modulated β-Band Oscillations Characterize Local Field Potential Activity in the Striatum of Awake Behaving Monkeys. J Neurosci 23:11741–11752.

Damier P, Hirsch EC, Agid Y, Graybiel AM (1999) The substantia nigra of the human. Brain 122:1437–1448.

Danielmeier C, Eichele T, Forstmann BU, Tittgemeyer M, Ullsperger M (2011) Posterior Medial Frontal Cortex Activity Predicts Post-Error Adaptations in Task-Related Visual and Motor Areas. J Neurosci 31:1780–1789.

Danielmeier C, Ullsperger M (2011) Post-error adjustments. Front Psychology 2:233.

De Lau LM, Breteler MM (2006) Epidemiology of Parkinson's disease. The Lancet Neurology 5:525–535.

DeLong MR (1990) Primate models of movement disorders of basal ganglia origin. Trends in Neurosciences 13:281–285.

Ding M, Bressler SL, Yang W, Liang H (2000) Short-window spectral analysis of cortical event-related potentials by adaptive multivariate autoregressive modeling: data preprocessing, model validation, and variability assessment. Biol Cybern 83:35–45.

Ding M, Chen Y, Bressler SL (2006) Granger Causality: Basis Theory and Application to Neuroscience. In: Handbook of Time Series Analysis: Recent Theoretical Developments and Applications (Schelter B, Winterhalder M, Timmer J, eds), pp 437–460. Berlin: Wiley-VCH. Available at: http://www.ccs.fau.edu/~bressler/pdf/HTSA06.pdf.

Eisenstein SA, Koller JM, Black KD, Campbell MC, Lugar HM, Ushe M, Tabbal SD, Karimi M, Hershey T, Perlmutter JS, Black KJ (2014) Functional anatomy of subthalamic nucleus stimulation in Parkinson disease. Ann Neurol 76:279–295.

Engel AK, Fries P (2010) Beta-band oscillations — signalling the status quo? Current Opinion in Neurobiology 20:156–165.

Eriksen BA, Eriksen CW (1974) Effects of noise letters upon the identification of a target letter in a nonsearch task. Perception & Psychophysics 16:143–149.

Eusebio A, Chen CC, Lu CS, Lee ST, Tsai CH, Limousin P, Hariz M, Brown P (2008) Effects of low-frequency stimulation of the subthalamic nucleus on movement in Parkinson's disease. Experimental Neurology 209:125–130.

Eusebio A, Thevathasan W, Gaynor LD, Pogosyan A, Bye E, Foltynie T, Zrinzo L, Ashkan K, Aziz T, Brown P (2011) Deep brain stimulation can suppress pathological synchronisation in parkinsonian patients. J Neurol Neurosurg Psychiatry 82:569–573.

Filion M, Tremblay L (1991) Abnormal spontaneous activity of globus pallidus neurons in monkeys with MPTP-induced parkinsonism. Brain Research 547:140–144.

Fitzgerald KD, Welsh RC, Gehring WJ, Abelson JL, Himle JA, Liberzon I, Taylor SF (2005) Error-related hyperactivity of the anterior cingulate cortex in obsessive-compulsive disorder. Biol Psychiatry 57:287–294.

Foltynie T, Hariz MI (2010) Surgical management of Parkinson's disease. Expert Review of Neurotherapeutics 10:903–914.

Frank MJ (2005) Dynamic Dopamine Modulation in the Basal Ganglia: A Neurocomputational Account of Cognitive Deficits in Medicated and Nonmedicated Parkinsonism. Journal of Cognitive Neuroscience 17:51–72.

Frank MJ (2006) Hold your horses: A dynamic computational role for the subthalamic nucleus in decision making. Neural Networks 19:1120–1136.

Frank MJ, Samanta J, Moustafa AA, Sherman SJ (2007) Hold Your Horses: Impulsivity, Deep Brain Stimulation, and Medication in Parkinsonism. Science 318:1309–1312.

Frank MJ, Seeberger LC, O'Reilly RC (2004) By Carrot or by Stick: Cognitive Reinforcement Learning in Parkinsonism. Science 306:1940–1943.

Fries P (2005) A mechanism for cognitive dynamics: neuronal communication through neuronal coherence. Trends in Cognitive Sciences 9:474–480.

Fumagalli M, Giannicola G, Rosa M, Marceglia S, Lucchiari C, Mrakic-Sposta S, Servello D, Pacchetti C, Porta M, Sassi M, Zangaglia R, Franzini A, Albanese A, Romito L, Piacentini S, Zago S, Pravettoni G, Barbieri S, Priori A (2011) Conflict-dependent dynamic of subthalamic nucleus oscillations during moral decisions. Soc Neurosci 6:243–256.

Gardner WA (1992) A unifying view of coherence in signal processing. Signal Processing 29:113–140.

Geller AS, Schlefer IK, Sederberg PB, Jacobs J, Kahana MJ (2007) PyEPL: a cross-platform experiment-programming library. Behav Res Methods 39:950–958.

Gerfen CR (1992) The Neostriatal Mosaic: Multiple Levels of Compartmental Organization in the Basal Ganglia. Annual Review of Neuroscience 15:285–320.

Gevins A, Smith ME, McEvoy L, Yu D (1997) High-resolution EEG mapping of cortical activation related to working memory: effects of task difficulty, type of processing, and practice. Cereb Cortex 7:374–385.

Geweke J (1982) Measurement of Linear Dependence and Feedback Between Multiple Time Series. Journal of the American Statistical Association 77:304–313.

Goldberg JH, Farries MA, Fee MS (2013) Basal ganglia output to the thalamus: still a paradox. Trends in Neurosciences 36:695–705.

Gold JI, Shadlen MN (2002) Banburismus and the Brain: Decoding the Relationship between Sensory Stimuli, Decisions, and Reward. Neuron 36:299–308.

Gold JI, Shadlen MN (2007) The neural basis of decision making. Annu Rev Neurosci 30:535–574.

Granger C (1969) Investigating Causal Relations by Econometric Models and Cross-spectral Methods. Econometrica 37:424–438.

Gratton G, Coles MGH, Donchin E (1992) Optimizing the use of information: Strategic control of activation of responses. Journal of Experimental Psychology: General 121:480–506.

Green N, Bogacz R, Huebl J, Beyer A-K, Kühn AA, Heekeren HR (2013) Reduction of Influence of Task Difficulty on Perceptual Decision Making by STN Deep Brain Stimulation. Current Biology 23:1681–1684.

Gross RE, Krack P, Rodriguez-Oroz MC, Rezai AR, Benabid A-L (2006) Electrophysiological mapping for the implantation of deep brain stimulators for Parkinson's disease and tremor. Mov Disord 21:S259–S283.

Hälbig TD, Tse W, Frisina PG, Baker BR, Hollander E, Shapiro H, Tagliati M, Koller WC, Olanow CW (2009) Subthalamic deep brain stimulation and impulse control in Parkinson's disease. Eur J Neurol 16:493–497.

Hammond C, Bergman H, Brown P (2007) Pathological synchronization in Parkinson's disease: networks, models and treatments. Trends in Neurosciences 30:357–364.

Hanslmayr S, Pastötter B, Bäuml K-H, Gruber S, Wimber M, Klimesch W (2007) The Electrophysiological Dynamics of Interference during the Stroop Task. Journal of Cognitive Neuroscience 20:215–225.

Hanson TL, Fuller AM, Lebedev MA, Turner DA, Nicolelis MAL (2012) Subcortical Neuronal Ensembles: An Analysis of Motor Task Association, Tremor, Oscillations, and Synchrony in Human Patients. J Neurosci 32:8620–8632.

Hashimoto T, Elder CM, Okun MS, Patrick SK, Vitek JL (2003) Stimulation of the Subthalamic Nucleus Changes the Firing Pattern of Pallidal Neurons. J Neurosci 23:1916–1923.

Hassler R (1938) Zur Pathologie der paralysis agitans und des postenzephalitischen Parkinsonismus. J Psychol Neurol 48:387–476.

Haynes WIA, Haber SN (2013) The Organization of Prefrontal-Subthalamic Inputs in Primates Provides an Anatomical Substrate for Both Functional Specificity and Integration: Implications for Basal Ganglia Models and Deep Brain Stimulation. J Neurosci 33:4804–4814.

Hazrati L-N, Parent A (1992) Convergence of subthalamic and striatal efferents at pallidal level in primates: an anterograde double-labeling study with biocytin and PHA-L. Brain Research 569:336–340.

Hernández-López S, Bargas J, Surmeier DJ, Reyes A, Galarraga E (1997) D1 Receptor Activation Enhances Evoked Discharge in Neostriatal Medium Spiny Neurons by Modulating an L-Type Ca2+ Conductance. J Neurosci 17:3334–3342.

Hernández-López S, Tkatch T, Perez-Garci E, Galarraga E, Bargas J, Hamm H, Surmeier DJ (2000) D2 Dopamine Receptors in Striatal Medium Spiny Neurons Reduce L-Type Ca2+ Currents and Excitability via a Novel PLCβ1–IP3–Calcineurin-Signaling Cascade. J Neurosci 20:8987–8995.

Hershey T, Campbell MC, Videen TO, Lugar HM, Weaver PM, Hartlein J, Karimi M, Tabbal SD, Perlmutter JS (2010) Mapping Go–No-Go performance within the subthalamic nucleus region. Brain 133:3625–3634.

Herz DM, Christensen MS, Bruggemann N, Hulme OJ, Ridderinkhof KR, Madsen KH, Siebner HR (2014) Motivational Tuning of Fronto-Subthalamic Connectivity Facilitates Control of Action Impulses. J Neurosci 34:3210–3217.

Herzog J, Reiff J, Krack P, Witt K, Schrader B, Müller D, Deuschl G (2003) Manic episode with psychotic symptoms induced by subthalamic nucleus stimulation in a patient with Parkinson's disease. Mov Disord 18:1382–1384.

Hornykiewicz O (2010) A brief history of levodopa. J Neurol 257:249–252.

Hughes AJ, Daniel SE, Kilford L, Lees AJ (1992) Accuracy of clinical diagnosis of idiopathic Parkinson's disease: a clinico-pathological study of 100 cases. J Neurol Neurosurg Psychiatr 55:181–184.

Ince E, Ciliax BJ, Levey AI (1997) Differential expression of D1 and D2 dopamine and m4 muscarinic acetylcholine receptor proteins in identified striatonigral neurons. Synapse 27:357–366.

Isoda M, Hikosaka O (2007) Switching from automatic to controlled action by monkey medial frontal cortex. Nat Neurosci 10:240–248.

Isoda M, Hikosaka O (2008) Role for Subthalamic Nucleus Neurons in Switching from Automatic to Controlled Eye Movement. J Neurosci 28:7209–7218.

Jahanshahi M, Ardouin CMA, Brown RG, Rothwell JC, Obeso J, Albanese A, Rodriguez-Oroz MC, Moro E, Benabid AL, Pollak P, Limousin-Dowsey P (2000) The impact of deep brain stimulation on executive function in Parkinson's disease. Brain 123:1142–1154.

Jahfari S, Verbruggen F, Frank MJ, Waldorp LJ, Colzato L, Ridderinkhof KR, Forstmann BU (2012) How Preparation Changes the Need for Top–Down Control of the Basal Ganglia When Inhibiting Premature Actions. J Neurosci 32:10870–10878.

Jentzsch I, Dudschig C (2009) Why do we slow down after an error? Mechanisms underlying the effects of posterror slowing. The Quarterly Journal of Experimental Psychology 62:209–218.

Joundi RA, Jenkinson N, Brittain J-S, Aziz TZ, Brown P (2012) Driving oscillatory activity in the human cortex enhances motor performance. Curr Biol 22:403–407.

Kamiński MJ, Blinowska KJ (1991) A new method of the description of the information flow in the brain structures. Biol Cybern 65:203–210.

Kerns JG, Cohen JD, MacDonald AW, Cho RY, Stenger VA, Carter CS (2004) Anterior Cingulate Conflict Monitoring and Adjustments in Control. Science 303:1023–1026.

Kitai ST, Sugimori M, Kocsis JD (1976) Excitatory nature of dopamine in the nigro-caudate pathway. Exp Brain Res 24:351–363.

Kordower JH, Olanow CW, Dodiya HB, Chu Y, Beach TG, Adler CH, Halliday GM, Bartus RT (2013) Disease duration and the integrity of the nigrostriatal system in Parkinson's disease. Brain 136:2419–2431.

Kravitz AV, Freeze BS, Parker PRL, Kay K, Thwin MT, Deisseroth K, Kreitzer AC (2010) Regulation of parkinsonian motor behaviours by optogenetic control of basal ganglia circuitry. Nature 466:622–626.

Kühn AA, Kupsch A, Schneider G-H, Brown P (2006) Reduction in subthalamic 8-35 Hz oscillatory activity correlates with clinical improvement in Parkinson's disease. European Journal of Neuroscience 23:1956–1960.

Kühn AA, Trottenberg T, Kivi A, Kupsch A, Schneider G-H, Brown P (2005) The relationship between local field potential and neuronal discharge in the subthalamic nucleus of patients with Parkinson's disease. Experimental Neurology 194:212–220.

Kühn AA, Williams D, Kupsch A, Limousin P, Hariz M, Schneider G-H, Yarrow K, Brown P (2004) Event-related beta desynchronization in human subthalamic nucleus correlates with motor performance. Brain 127:735–746.

Lachaux J-P, Lutz A, Rudrauf D, Cosmelli D, Le Van Quyen M, Martinerie J, Varela F (2002) Estimating the time-course of coherence between single-trial brain signals: an introduction to wavelet coherence. Neurophysiologie Clinique/Clinical Neurophysiology 32:157–174.

Laitinen LV, Bergenheim AT, Hariz MI (1992) Leksell's posteroventral pallidotomy in the treatment of Parkinson's disease. Journal of Neurosurgery 76:53–61.

Langston JW, Ballard P, Tetrud JW, Irwin I (1983) Chronic Parkinsonism in humans due to a product of meperidine-analog synthesis. Science 219:979–980.

Lardeux S, Pernaud R, Paleressompoulle D, Baunez C (2009) Beyond the reward pathway: coding reward magnitude and error in the rat subthalamic nucleus. J Neurophysiol 102:2526–2537.

Lesage S, Brice A (2009) Parkinson's disease: from monogenic forms to genetic susceptibility factors. Hum Mol Genet 18:R48–R59.

Leventhal DK, Gage GJ, Schmidt R, Pettibone JR, Case AC, Berke JD (2012) Basal Ganglia Beta Oscillations Accompany Cue Utilization. Neuron 73:523–536.

Levy R, Ashby P, Hutchison WD, Lang AE, Lozano AM, Dostrovsky JO (2002) Dependence of subthalamic nucleus oscillations on movement and dopamine in Parkinson's disease. Brain 125:1196–1209.

Levy R, Hutchison WD, Lozano AM, Dostrovsky JO (2000) High-frequency Synchronization of Neuronal Activity in the Subthalamic Nucleus of Parkinsonian Patients with Limb Tremor. J Neurosci 20:7766–7775.

Limousin P, Krack P, Pollak P, Benazzouz A, Ardouin C, Hoffmann D, Benabid A-L (1998) Electrical Stimulation of the Subthalamic Nucleus in Advanced Parkinson's Disease. New England Journal of Medicine 339:1105–1111.

Limousin P, Pollak P, Benazzouz A, Hoffmann D, Le Bas JF, Broussolle E, Perret JE, Benabid AL (1995) Effect of parkinsonian signs and symptoms of bilateral subthalamic nucleus stimulation. Lancet 345:91–95.

Lipszyc J, Schachar R (2010) Inhibitory control and psychopathology: A meta-analysis of studies using the stop signal task. Journal of the International Neuropsychological Society 16:1064–1076.

Little S, Pogosyan A, Neal S, Zavala B, Zrinzo L, Hariz M, Foltynie T, Limousin P, Ashkan K, FitzGerald J, Green AL, Aziz TZ, Brown P (2013) Adaptive deep brain stimulation in advanced Parkinson disease. Ann Neurol 74:449–457.

Logan GD, Cowan WB, Davis KA (1984) On the ability to inhibit simple and choice reaction time responses: A model and a method. Journal of Experimental Psychology: Human Perception and Performance 10:276–291.

Lu C, Bharmal A, Suchowersky O (2006) Gambling and parkinson disease. Arch Neurol 63:298–298.

Magariños-Ascone CM, Figueiras-Mendez R, Riva-Meana C, Córdoba-Fernández A (2000) Subthalamic neuron activity related to tremor and movement in Parkinson's disease. European Journal of Neuroscience 12:2597–2607.

Mallet L, Schüpbach M, N'Diaye K, Remy P, Bardinet E, Czernecki V, Welter M-L, Pelissolo A, Ruberg M, Agid Y, Yelnik J (2007) Stimulation of subterritories of the subthalamic nucleus reveals its role in the integration of the emotional and motor aspects of behavior. PNAS 104:10661–10666.

Mandat TS, Hurwitz T, Honey CR (2006) Hypomania as an adverse effect of subthalamic nucleus stimulation: report of two cases. Acta Neurochir (Wien) 148:895–898.

Mann JM, Foote KD, Garvan CW, Fernandez HH, Jacobson CE, Rodriguez RL, Haq IU, Siddiqui MS, Malaty IA, Morishita T, Hass CJ, Okun MS (2009) Brain penetration effects of microelectrodes and DBS leads in STN or GPi. J Neurol Neurosurg Psychiatr 80:794–797.

Maris E, Oostenveld R (2007) Nonparametric statistical testing of EEG- and MEG-data. J Neurosci Methods 164:177–190.

Marsden CD (1990) Parkinson's disease. The Lancet 335:948–949.

Marsden CD, Obeso JA (1994) The functions of the basal ganglia and the paradox of stereotaxic surgery in Parkinson's disease. Brain 117:877–897.

McIntyre CC, Hahn PJ (2010) Network Perspectives on the Mechanisms of Deep Brain Stimulation. Neurobiol Dis 38:329–337.

Miller WC, DeLong MR (1987) Altered Tonic Activity of Neurons in the Globus Pallidus and Subthalamic Nucleus in the Primate MPTP Model of Parkinsonism. In: The Basal Ganglia II (Carpenter MB, Jayaraman A, eds), pp 415–427 Advances in Behavioral Biology. Springer US. Available at: http://link.springer.com/chapter/10.1007/978-1-4684-5347-8_29 [Accessed December 1, 2014].

Minhas P, Bansal V, Patel J, Ho JS, Diaz J, Datta A, Bikson M (2010) Electrodes for high-definition transcutaneous DC stimulation for applications in drug delivery and electrotherapy, including tDCS. J Neurosci Methods 190:188–197.

Mirabella G, Iaconelli S, Romanelli P, Modugno N, Lena F, Manfredi M, Cantore G (2012) Deep Brain Stimulation of Subthalamic Nuclei Affects Arm Response Inhibition In Parkinson's Patients. Cereb Cortex 22:1124–1132.

Monakow KH, Akert K, Künzle H (1978) Projections of the precentral motor cortex and other cortical areas of the frontal lobe to the subthalamic nucleus in the monkey. Exp Brain Res 33:395–403.

Moran A, Stein E, Tischler H, Bar-Gad I (2012) Decoupling neuronal oscillations during subthalamic nucleus stimulation in the parkinsonian primate. Neurobiol Dis 45:583–590.

Nachev P (2011) The blind executive. NeuroImage 57:312–313.

Nambu A, Takada M, Inase M, Tokuno H (1996) Dual somatotopical representations in the primate subthalamic nucleus: evidence for ordered but reversed body-map transformations from the primary motor cortex and the supplementary motor area. J Neurosci 16:2671–2683.

Nambu A, Tokuno H, Hamada I, Kita H, Imanishi M, Akazawa T, Ikeuchi Y, Hasegawa N (2000) Excitatory Cortical Inputs to Pallidal Neurons Via the Subthalamic Nucleus in the Monkey. Journal of Neurophysiology 84:289–300.

Nambu A, Tokuno H, Takada M (2002) Functional significance of the cortico–subthalamo–pallidal "hyperdirect" pathway. Neuroscience Research 43:111–117.

Narabayashi H, Okuma T (1953) Procaine-oil blocking of the globus pallidus for the treatment of rigidity and tremor of parkinsonism. Proc Jpn Acad 29:134–137.

Narayanan NS, Cavanagh JF, Frank MJ, Laubach M (2013) Common medial frontal mechanisms of adaptive control in humans and rodents. Nat Neurosci 16:1888–1895.

Narayanan NS, Laubach M (2009) Delay Activity in Rodent Frontal Cortex During a Simple Reaction Time Task. Journal of Neurophysiology 101:2859–2871.

Neubert F-X, Mars RB, Buch ER, Olivier E, Rushworth MFS (2010) Cortical and subcortical interactions during action reprogramming and their related white matter pathways. Proc Natl Acad Sci USA 107:13240–13245.

Obeso I, Wilkinson L, Casabona E, Speekenbrink M, Bringas ML, Álvarez M, Álvarez L, Pavón N, Rodríguez-Oroz MC, Macías R, Obeso JA, Jahanshahi M (2014) The subthalamic nucleus and inhibitory control: impact of subthalamotomy in Parkinson's disease. Brain 137:1470–1480.

Obeso JA, Rodriguez MC, DeLong MR (1997) Basal ganglia pathophysiology. A critical review. Adv Neurol 74:3–18.

O'Connell RG, Dockree PM, Kelly SP (2012) A supramodal accumulation-to-bound signal that determines perceptual decisions in humans. Nature Neuroscience 15:1729–1735.

Oehrn CR, Hanslmayr S, Fell J, Deuker L, Kremers NA, Lam ATD, Elger CE, Axmacher N (2014) Neural Communication Patterns Underlying Conflict Detection, Resolution, and Adaptation. J Neurosci 34:10438–10452.

Oswal A, Litvak V, Sauleau P, Brown P (2012) Beta Reactivity, Prospective Facilitation of Executive Processing, and Its Dependence on Dopaminergic Therapy in Parkinson's Disease. J Neurosci 32:9909–9916.

Peirce JW (2007) PsychoPy—Psychophysics software in Python. Journal of Neuroscience Methods 162:8–13.

Pesaran B (2008) Short course III, presented at 2008 Society for Neuroscience Annual Meeting Neural signal processing: quantitative analysis of neural activity (Mitra P, ed). Washington DC: Society for Neuroscience.

Pizzagalli DA, Oakes TR, Davidson RJ (2003) Coupling of theta activity and glucose metabolism in the human rostral anterior cingulate cortex: an EEG/PET study of normal and depressed subjects. Psychophysiology 40:939–949.

Pogosyan A, Gaynor LD, Eusebio A, Brown P (2009) Boosting cortical activity at Beta-band frequencies slows movement in humans. Curr Biol 19:1637–1641.

Pollo C, Kaelin-Lang A, Oertel MF, Stieglitz L, Taub E, Fuhr P, Lozano AM, Raabe A, Schüpbach M (2014) Directional deep brain stimulation: an intraoperative double-blind pilot study. Brain 137:2015–2026.

Rabbitt PM (1966) Errors and error correction in choice-response tasks. J Exp Psychol 71:264–272.

Rajput AH, Sitte HH, Rajput A, Fenton ME, Pifl C, Hornykiewicz O (2008) Globus pallidus dopamine and Parkinson motor subtypes Clinical and brain biochemical correlation. Neurology 70:1403–1410.

Ratcliff R, Frank MJ (2012) Reinforcement-based decision making in corticostriatal circuits: mutual constraints by neurocomputational and diffusion models. Neural Comput 24:1186–1229.

Ratcliff R, McKoon G (2008) The diffusion decision model: theory and data for two-choice decision tasks. Neural Comput 20:873–922.

Ray NJ, Brittain J-S, Holland P, Joundi RA, Stein JF, Aziz TZ, Jenkinson N (2012) The role of the subthalamic nucleus in response inhibition: Evidence from local field potential recordings in the human subthalamic nucleus. NeuroImage 60:271–278.

Ray NJ, Jenkinson N, Brittain J, Holland P, Joint C, Nandi D, Bain PG, Yousif N, Green A, Stein JS, Aziz TZ (2009) The role of the subthalamic nucleus in response inhibition: Evidence from deep brain stimulation for Parkinson's disease. Neuropsychologia 47:2828–2834.

Ray NJ, Jenkinson N, Wang S, Holland P, Brittain JS, Joint C, Stein JF, Aziz T (2008) Local field potential beta activity in the subthalamic nucleus of patients with Parkinson's

disease is associated with improvements in bradykinesia after dopamine and deep brain stimulation. Experimental Neurology 213:108–113.

Ridderinkhof KR, Ullsperger M, Crone EA, Nieuwenhuis S (2004) The Role of the Medial Frontal Cortex in Cognitive Control. Science 306:443–447.

Rodriguez-Oroz MC, López-Azcárate J, Garcia-Garcia D, Alegre M, Toledo J, Valencia M, Guridi J, Artieda J, Obeso JA (2011) Involvement of the subthalamic nucleus in impulse control disorders associated with Parkinson's disease. Brain 134:36–49.

Rodriguez-Oroz MC, Rodriguez M, Guridi J, Mewes K, Chockkman V, Vitek J, DeLong MR, Obeso JA (2001) The subthalamic nucleus in Parkinson's disease: somatotopic organization and physiological characteristics. Brain 124:1777–1790.

Romito LM, Raja M, Daniele A, Contarino MF, Bentivoglio AR, Barbier A, Scerrati M, Albanese A (2002) Transient mania with hypersexuality after surgery for high frequency stimulation of the subthalamic nucleus in Parkinson's disease. Mov Disord 17:1371–1374.

Rosa M, Fumagalli M, Giannicola G, Marceglia S, Lucchiari C, Servello D, Franzini A, Pacchetti C, Romito L, Albanese A, Porta M, Pravettoni G, Priori A (2013) Pathological gambling in Parkinson's disease: Subthalamic oscillations during economics decisions. Mov Disord 28:1644–1652.

Rosin B, Slovik M, Mitelman R, Rivlin-Etzion M, Haber SN, Israel Z, Vaadia E, Bergman H (2011) Closed-Loop Deep Brain Stimulation Is Superior in Ameliorating Parkinsonism. Neuron 72:370–384.

Ruiz MH, Huebl J, Schönecker T, Kupsch A, Yarrow K, Krauss JK, Schneider G-H, Kühn AA (2013) Involvement of Human Internal Globus Pallidus in the Early Modulation of Cortical Error-Related Activity. Cereb Cortex:bht002.

Ruiz MH, Strübing F, Jabusch H-C, Altenmüller E (2011) EEG oscillatory patterns are associated with error prediction during music performance and are altered in musician's dystonia. NeuroImage 55:1791–1803.

Rustamov N, Rodriguez-Raecke R, Timm L, Agrawal D, Dressler D, Schrader C, Tacik P, Wegner F, Dengler R, Wittfoth M, Kopp B (2013) Absence of congruency sequence effects reveals neurocognitive inflexibility in Parkinson's disease. Neuropsychologia 51:2976–2987.

Satoh T, Nakai S, Sato T, Kimura M (2003) Correlated Coding of Motivation and Outcome of Decision by Dopamine Neurons. J Neurosci 23:9913–9923.

Scherbaum S, Dshemuchadse M (2013) Higher response time increases theta energy, conflict increases response time. Clin Neurophysiol 124:1477–1479.

Schlögl A, Supp G (2006) Analyzing event-related EEG data with multivariate autoregressive parameters. Prog Brain Res 159:135–147.

Schmidt R, Leventhal DK, Mallet N, Chen F, Berke JD (2013) Canceling actions involves a race between basal ganglia pathways. Nat Neurosci 16:1118–1124.

Schroeder U, Kuehler A, Haslinger B, Erhard P, Fogel W, Tronnier VM, Lange KW, Boecker H, Ceballos-Baumann AO (2002) Subthalamic nucleus stimulation affects striato-anterior cingulate cortex circuit in a response conflict task: a PET study. Brain 125:1995–2004.

Schultz W (2002) Getting Formal with Dopamine and Reward. Neuron 36:241–263.

Schultz W, Dayan P, Montague PR (1997) A Neural Substrate of Prediction and Reward. Science 275:1593–1599.

Schwarz G (1978) Estimating the Dimension of a Model. Ann Statist 6:461–464.

Seth AK (2010) A MATLAB toolbox for Granger causal connectivity analysis. Journal of Neuroscience Methods 186:262–273.

Sheth SA, Mian MK, Patel SR, Asaad WF, Williams ZM, Dougherty DD, Bush G, Eskandar EN (2012) Human dorsal anterior cingulate cortex neurons mediate ongoing behavioural adaptation. Nature 488:218–221.

Siegert S, Herrojo Ruiz M, Brücke C, Huebl J, Schneider G-H, Ullsperger M, Kühn AA (2014) Error signals in the subthalamic nucleus are related to post-error slowing in patients with Parkinson's disease. Cortex.

Silberstein P, Pogosyan A, Kühn AA, Hotton G, Tisch S, Kupsch A, Dowsey-Limousin P, Hariz MI, Brown P (2005) Cortico-cortical coupling in Parkinson's disease and its modulation by therapy. Brain 128:1277–1291.

Simon JR (1969) Reactions toward the source of stimulation. J Exp Psychol 81:174–176.

Smeding HMM, Goudriaan AE, Foncke EMJ, Schuurman PR, Speelman JD, Schmand B (2007) Pathological gambling after bilateral subthalamic nucleus stimulation in Parkinson disease. J Neurol Neurosurg Psychiatr 78:517–519.

Smith Y, Bevan MD, Shink E, Bolam JP (1998) Microcircuitry of the direct and indirect pathways of the basal ganglia. Neuroscience 86:353–387.

Spinks RL, Kraskov A, Brochier T, Umilta MA, Lemon RN (2008) Selectivity for Grasp in Local Field Potential and Single Neuron Activity Recorded Simultaneously from M1 and F5 in the Awake Macaque Monkey. J Neurosci 28:10961–10971.

Stefani A, Fedele E, Galati S, Pepicelli O, Frasca S, Pierantozzi M, Peppe A, Brusa L, Orlacchio A, Hainsworth AH, Gattoni G, Stanzione P, Bernardi G, Raiteri M, Mazzone P (2005) Subthalamic stimulation activates internal pallidus: Evidence from cGMP microdialysis in PD patients. Ann Neurol 57:448–452.

Stroop RJ (1935) Studies of interference in serial verbal reactions. Journal of Experimental Psychology 18:643–662.

Svennilson E, Torvik A, Lowe R, Leksell L (1960) Treatment of parkinsonism by stereotatic thermolesions in the pallidal region. A clinical evaluation of 81 cases. Acta Psychiatr Scand 35:358–377.

Swann NC, Cai W, Conner CR, Pieters TA, Claffey MP, George JS, Aron AR, Tandon N (2012) Roles for the pre-supplementary motor area and the right inferior frontal gyrus in stopping action: electrophysiological responses and functional and structural connectivity. Neuroimage 59:2860–2870.

Swann N, Poizner H, Houser M, Gould S, Greenhouse I, Cai W, Strunk J, George J, Aron AR (2011) Deep Brain Stimulation of the Subthalamic Nucleus Alters the Cortical Profile of Response Inhibition in the Beta Frequency Band: A Scalp EEG Study in Parkinson's Disease. J Neurosci 31:5721–5729.

Swann N, Tandon N, Canolty R, Ellmore TM, McEvoy LK, Dreyer S, DiSano M, Aron AR (2009) Intracranial EEG Reveals a Time- and Frequency-Specific Role for the Right Inferior Frontal Gyrus and Primary Motor Cortex in Stopping Initiated Responses. J Neurosci 29:12675–12685.

Tanner CM, Ottman R, Goldman SM, et al (1999) Parkinson disease in twins: An etiologic study. JAMA 281:341–346.

Tokuno H, Nambu A (2000) Organization of Nonprimary Motor Cortical Inputs on Pyramidal and Nonpyramidal Tract Neurons of Primary Motor Cortex: An Electrophysiological Study in the Macaque Monkey. Cereb Cortex 10:58–68.

Van den Wildenberg WPM, van Boxtel GJM, van der Molen MW, Bosch DA, Speelman JD, Brunia CHM (2006) Stimulation of the Subthalamic Region Facilitates the Selection and Inhibition of Motor Responses in Parkinson's Disease. Journal of Cognitive Neuroscience 18:626–636.

Van Driel J, Ridderinkhof KR, Cohen MX (2012) Not All Errors Are Alike: Theta and Alpha EEG Dynamics Relate to Differences in Error-Processing Dynamics. J Neurosci 32:16795–16806.

Van Meel CS, Heslenfeld DJ, Oosterlaan J, Sergeant JA (2007) Adaptive control deficits in attention-deficit/hyperactivity disorder (ADHD): the role of error processing. Psychiatry Res 151:211–220.

Vinck M, van Wingerden M, Womelsdorf T, Fries P, Pennartz CMA (2010) The pairwise phase consistency: A bias-free measure of rhythmic neuronal synchronization. NeuroImage 51:112–122.

Voon V, Fox SH (2007) MEdication-related impulse control and repetitive behaviors in parkinson disease. Arch Neurol 64:1089–1096.

Wang C, Ulbert I, Schomer DL, Marinkovic K, Halgren E (2005) Responses of Human Anterior Cingulate Cortex Microdomains to Error Detection, Conflict Monitoring, Stimulus-Response Mapping, Familiarity, and Orienting. J Neurosci 25:604–613.

Weinberger M, Mahant N, Hutchison WD, Lozano AM, Moro E, Hodaie M, Lang AE, Dostrovsky JO (2006) Beta oscillatory activity in the subthalamic nucleus and its relation to dopaminergic response in Parkinson's disease. J Neurophysiol 96:3248–3256.

Whitmer D, de Solages C, Hill B, Yu H, Henderson JM, Bronte-Stewart H (2012) High frequency deep brain stimulation attenuates subthalamic and cortical rhythms in Parkinson's disease. Front Hum Neurosci 6:155.

Wichmann T, Bergman H, DeLong MR (1994) The primate subthalamic nucleus. III. Changes in motor behavior and neuronal activity in the internal pallidum induced by subthalamic inactivation in the MPTP model of parkinsonism. J Neurophysiol 72:521–530.

Wichmann T, DeLong MR, Guridi J, Obeso JA (2011) Milestones in research on the pathophysiology of Parkinson's disease. Mov Disord 26:1032–1041.

Wichmann T, Soares J (2006) Neuronal Firing Before and After Burst Discharges in the Monkey Basal Ganglia Is Predictably Patterned in the Normal State and Altered in Parkinsonism. Journal of Neurophysiology 95:2120–2133.

Wiecki TV, Frank MJ (2013) A computational model of inhibitory control in frontal cortex and basal ganglia. Psychol Rev 120:329–355.

Williams D, Tijssen M, Bruggen G van, Bosch A, Insola A, Lazzaro VD, Mazzone P, Oliviero A, Quartarone A, Speelman H, Brown P (2002) Dopamine-dependent changes in the functional connectivity between basal ganglia and cerebral cortex in humans. Brain 125:1558–1569.

Williams ZM, Neimat JS, Cosgrove GR, Eskandar EN (2005) Timing and direction selectivity of subthalamic and pallidal neurons in patients with Parkinson disease. Exp Brain Res 162:407–416.

Windels F, Bruet N, Poupard A, Urbain N, Chouvet G, Feuerstein C, Savasta M (2000) Effects of high frequency stimulation of subthalamic nucleus on extracellular glutamate and GABA in substantia nigra and globus pallidus in the normal rat. European Journal of Neuroscience 12:4141–4146.

Wise SP (2008) Forward frontal fields: phylogeny and fundamental function. Trends in Neurosciences 31:599–608.

Witt K, Pulkowski U, Herzog J, et al (2004) Deep brain stimulation of the subthalamic nucleus improves cognitive flexibility but impairs response inhibition in parkinson disease. Arch Neurol 61:697–700.

Wylie SA, Ridderinkhof KR, Elias WJ, Frysinger RC, Bashore TR, Downs KE, Wouwe NC van, Wildenberg WPM van den (2010) Subthalamic nucleus stimulation influences expression and suppression of impulsive behaviour in Parkinson's disease. Brain 133:3611–3624.

Yao SC, Hart AD, Terzella MJ (2013) An evidence-based osteopathic approach to Parkinson disease. Osteopathic Family Physician 5:96–101.

Yeung N, Botvinick MM, Cohen JD (2004) The neural basis of error detection: conflict monitoring and the error-related negativity. Psychol Rev 111:931–959.

Yeung N, Cohen JD, Botvinick MM (2011) Errors of interpretation and modeling: a reply to Grinband et al. Neuroimage 57:316–319.

Yordanova J, Falkenstein M, Hohnsbein J, Kolev V (2004) Parallel systems of error processing in the brain. NeuroImage 22:590–602.

Zaghloul KA, Blanco JA, Weidemann CT, McGill K, Jaggi JL, Baltuch GH, Kahana MJ (2009) Human Substantia Nigra Neurons Encode Unexpected Financial Rewards. Science 323:1496–1499.

Zaghloul KA, Weidemann CT, Lega BC, Jaggi JL, Baltuch GH, Kahana MJ (2012) Neuronal Activity in the Human Subthalamic Nucleus Encodes Decision Conflict during Action Selection. J Neurosci 32:2453–2460.

Zavala BA, Tan H, Little S, Ashkan K, Hariz M, Foltynie T, Zrinzo L, Zaghloul KA, Brown P (2014) Midline Frontal Cortex Low-Frequency Activity Drives Subthalamic Nucleus Oscillations during Conflict. J Neurosci 34:7322–7333.

Zavala B, Brittain J-S, Jenkinson N, Ashkan K, Foltynie T, Limousin P, Zrinzo L, Green AL, Aziz T, Zaghloul K, Brown P (2013) Subthalamic Nucleus Local Field Potential Activity during the Eriksen Flanker Task Reveals a Novel Role for Theta Phase during Conflict Monitoring. J Neurosci 33:14758–14766.