Quantifying the subjective experience of initiating and monitoring actions, in health and disease.

Steven Di Costa

Institute of Cognitive Neuroscience

University College London

PhD Thesis

2017

Declaration

I, Steven Di Costa, 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.

Steven Di Costa

Abstract

Human movement is a ubiquitous behaviour that is performed to achieve goals by affecting changes in the outside world. The creation of the manmade environment has depended on our ability to plan, execute and monitor movements. This thesis will investigate the subjective experience and underlying brain processes that accompany voluntary and outcome-directed actions.

Experiments in this work will make use of an established paradigm to quantify and compare participants' perceptions of their own actions. Specifically, through the use of a rotating clock, participants report when they perceived action-relevant events to occur, such as the intention to act, the action itself, or the outcome produced by the action.

The experience of producing an action, even in the absence of a desired outcome, varies among populations. This thesis will compare the experience of volition in patients with Parkinson's disease and in healthy controls. It will be seen that manipulation of dopamine availability, either through medication or deep brain stimulation, influences the subjective experience of making an action. Using EEG, it will be seen that differences in the experience of volition are associated with the time-course of activation in brain areas typically related to movement initiation.

The sense of agency, a feeling of control of over one's own actions and their outcomes, will also be investigated by combining the paradigm described above with an established decision-making task. Results from a series of experiments will reveal a sequential effect in which an implicit measure of the sense of agency (intentional binding) is stronger following actions that produce negative outcomes. Finally, this novel effect, which appears to be related to learning processes, will be subjected to meta-analysis.

The results obtained in this work will inform further research into volition, motor-control and movement-related pathology.

Acknowledgments

I am immensely grateful to my supervisor, Prof. Patrick Haggard, for his guidance and patience throughout my PhD. It has been a privilege to work with him and to benefit from his knowledge, insights and good humour. His leadership of the Action & Body lab is a rare mix of brilliant instruction and unwavering kindness. As a result, I have spent the past four years working with inspiring colleagues and excellent friends.

I would particularly like to thank Nima Khalighinejad, Brianna Beck and Lieke de Boer for their companionship, advice and support over the years.

Chapters 2 and 3 would not have been possible without the collaboration and hard work of Christos Ganos, Ewgenia Barow, Christian Moll and Allessandro Gulberti, in Hamburg.

For Chapters 4 and 5, I wish to thank Héloïse Théro for her collaboration, advice and data collection. Her insights had a profound influence on the interpretation of data in these chapters. I would also like to thank Valérian Chambon for his thoughtful suggestions and support. The work described in Chapter 4 has partly been published as "Di Costa, S., Théro, H., Chambon, V., & Haggard, P. (2017). Try and try again: Post-error boost of an implicit measure of agency. *The Quarterly Journal of Experimental Psychology, 1*-28."

For Chapter 7, I would like to thank Francesca Morgante and Lucia Ricciardi in Messina for inviting me to work in their clinic and for their enthusiastic collaboration.

I wish to thank my family in Australia who, despite the distance, have been a constant source of love and support.

Finally, I would like to thank my boyfriend, Dean Garland, for the tea, and everything that comes with it.

Contents

List of figures

List of tables

Chapter 1. Introduction

1.1. Perpetual motion.

Human movement is constant. Like all ambulant organisms, we do not stop moving, even when at rest. Consider the innumerable processes required for you to be able to sit and read this document. You are engaged in keeping your body upright and balanced. You may shift yourself for comfort. You may scratch an itch. You are semi-automatically moving your eyes across the pages, which you are deliberately turning with your hands. You are blinking. You are breathing.

Even within this limited context, the extent and variety of movements is staggering. Many of these movements, such as breathing, blinking and sneezing, will occur without you being aware of any specific decision to make them happen. Some are reflexes. Other movements will occur seemingly as a result of an intervention or decision on 'your' part. When you turn the pages, glance at your wristwatch or reach for a mug of coffee, you are performing what are generally considered voluntary actions. In purely mechanical terms, there is no reason why these different classes of movement should be phenomenologically distinct. The muscular activity for voluntary and involuntary movements can be identical, yet we experience them very differently because of the types of stimulation that give rise to them and the underlying brain processes that accompany them.

The subjective experiences that we associate with voluntary actions will be the primary focus of this thesis.

1.2. Voluntary actions

1.2.1 Definitions and focus

An empirical study of voluntary action requires careful definitions. Unlike for temperature and weight, we have no standardized measurements for phenomena like intent, will or control. Furthermore, everyday descriptions of these experiences are shaped by societal factors like religion, law and language.

For the purposes of this thesis, two specific aspects of voluntary action will be considered separately: **volition,** the intention to act, which precedes the physical movement (see section 1.3.), and **agency**, the feeling of control over one's actions and their resulting effects (see section 1.4.).

1.2.2. What are voluntary actions for?

Voluntary actions come at a temporal and energetic cost. Deliberation and goal-directed planning takes time, even for simple gestures. Reflexes and habitual actions are faster, but are also inflexible and sometimes inappropriate to specific contexts (Keremati, Dezfouli, & Piray, 2011). The capacity to delay or alter certain kinds of actions may therefore prove beneficial to organisms that operate within a complex environment.

Voluntary actions may also help us to develop a sophisticated understanding of the environments we inhabit and how to achieve goals and obtain rewards in shifting contexts. Indeed, many paradigms in experimental psychology assume and rely on the voluntary action system to investigate learning in unstable choice environments (see for example Rolls, 2000; Cools, Clark Owen, & Robbins, 2002). On one view, the

voluntariness of an action simply refers to the relatively sophisticated way it is adapted to context and environment (Schüür & Haggard, 2011).

Perhaps one of the clearest indicators of the importance of voluntary action is the psychopathology that often accompanies its absence or disruption.

1.2.3. Absent or disordered voluntary actions

Movement paralysis is invariably distressing. The inability to move represents the inability to pursue goals effectively and efficiently. Particular discomfort may also arise from the disruption of movements that are used to communicate, such as facial expression and gesturing (Coulson, O'Dwyer, Adams, & Croxson, 2004). This type of disruption may be equally disturbing for observers and interlocutors who rely on the perception of such movements to interpret the thoughts and feelings of others. The social interactions of schizophrenic patients with flat affect, for example, are typically impaired relative to healthy controls (Gur et al., 2006).

Disturbances of the voluntary movement system are a key symptom of several pathologies, including Parkinson's disease, Huntington's chorea, dystonia and Tourette's syndrome. In Tourette's syndrome, for example, involuntary movements (tics) may be chronic and impair everyday function. In such cases medication is often prescribed. Yet even mild tics require family support, environment restructuring and on-going therapy (Kurlan, 1997). The impact of impaired voluntary movement is a critical contributor to recommended measures of health-related quality of life, for example in Parkinson's disease (Martinez-Martin, Rodriguez-Blazquez, Kurtis, & Chaudhuri, 2011) and recovery of voluntary movement is considered a priority

treatment in recovery from brain damage, for example in stroke patients (Turton & Fraser, 1986).

Disturbances in voluntary movement may also arise from the *presence* of an involuntary movement, rather than *loss* of volition. The Kohnstamm phenomenon, for example, describes the feeling of 'floating arms' after prolonged isometric engagement of the deltoid muscle against a resistant surface (Kohnstamm, 1915). The effect is short-lived (less than a minute), not functionally disruptive and therefore not normally experienced as distressing. However, the phenomenon is sufficiently unfamiliar to be considered a curious departure from our experience of voluntary movement and has been a subject of investigation for more than 100 years (De Havas, Ghosh, Gomi, & Haggard, 2016). Examples such as this strongly suggest that that our default experience of movement incorporates voluntariness. The absence of volition or the presence of involuntary movements is generally experienced as strange and considered pathological.

1.2.4. Investigating voluntary movements – aims and motivation

Given the importance of voluntary movement, it is perhaps unsurprising that it has long been a topic of scientific enquiry. These investigations have done more than just satisfy intellectual curiosity. Understanding how voluntary movements arise has been critical to the development of movement disorder therapies and pharmacological interventions. These interventions enable people to achieve objectives, when their ability to do so has become compromised.

In very different circumstances, the study of movement may also be useful in protecting those who lack control over their actions. Legal systems make an important distinction between criminal actions that were performed with or without volition on the part of

the accused. When commission of the act itself is not contested, sentencing typically depends on whether or not the accused knew what they were doing, and understood the consequences of their actions (*mens rea,* in legal terms). If this cognitive 'appreciation' of the crime is absent, a defence of insanity may be considered (Ogloff, 1991; Garrison, 1998). Regrettably, it has been shown in mock juror group studies that understanding the extent and application of formal guidelines (e.g. the M'Naughton rules or the ALI rule) differs significantly between individuals (Ogloff, 1991). In cases where subjective judgements are required, individual differences in conceptual understanding may influence decisions. The use of objective criteria would therefore be preferred, particularly in high-stakes contexts, such as criminal law. A more comprehensive understanding of the quantifiable aspects of voluntary movements is therefore desirable.

1.2.5. Investigating voluntary movements – methodology

The study of voluntary movement was historically associated with the description and treatment of movement disorders. This tradition dates back at least to Galen in his work *'de Tremore'* [trad: on tremors] (Louis, 2000). In even earlier texts, frequent mention of possession by malicious spirits, spasms, fits and seizures suggests that disturbances in voluntary action were noted and studied at least as far back as there are written records.

In the absence of disordered movement, empirical investigation of the mechanisms underlying voluntary actions was for many centuries undermined by religious beliefs in a soul, or similar concepts. Widespread dualism presupposed that a non-physical aspect of the self made decisions about actions and caused them to manifest. However, the operation of the central nervous system, including the motor system, has been the subject of speculation and experimentation since antiquity (Finger, 2001). It is the

subjective experience accompanying voluntary movements that was considered a subject for philosophers and theologians until the emergence of a scientific psychology.

William James, in *Principles* (James, 1890) makes a very clear distinction between reflexes or automatic movements, which are "unforeseen" by the agent, and actions that are "desired and intended" (p. 486). James places a particular emphasis on the prediction of the sensory consequences of one's action (with particular reference to visual input). This approach may be considered a precursor to the internal model principle (Francis & Wonham, 1976), according to which the central nervous system (CNS) makes a prediction about the result of a voluntary action, which is compared to the sensory result to provide feedback to the system. In this sense, as early as the 19th century there were attempts to understand the experience of performing an action and then monitoring its consequences.

These early conceptual frameworks were later combined with the emerging psychophysical methods developed to quantify and measure aspects of voluntary action such as reaction time (Galton, 1889). Experimental work would continue in the psychophysical tradition to examine related variables such as force exertion (Stevens & Mack, 1959), and time perception (Libet et al., 1983).

Attempts to better describe and understand the mechanisms underlying voluntary action have also been assisted by improved technology for studying the structure and function of the brain. Transcranial magnetic stimulation (TMS) has facilitated the comparison of voluntary and involuntary movements in a controlled setting (see for example Haggard, Clark, & Kalogeras, 2002). Electroencephalography (EEG) has enabled a more precise examination of the temporal components of brain activity preceding a voluntary action (Libet, et al., 1983), while experiments using fMRI have

provided more fine-grained spatial descriptions of neuronal activity in a non-surgical setting (Lotze et al., 1999). Specific experimental paradigms relevant to volition and agency are described below.

1.2.6. Aims of this thesis

Despite much progress, scientific understanding of the subjective experience that accompanies voluntary actions is far from comprehensive.

Patient groups with disordered movement may experience a markedly different neural context in which their actions occur. Specific brain sites may be damaged (as in stroke), or neurotransmission may be impaired (as in Parkinsons's disease). Yet the effect these altered contexts have on the subjective experience of voluntary actions is not well defined beyond what patients themselves report.

Furthermore, many of the established paradigms for studying voluntary action rely on repetitive movements with very little incentive or meaning attached to them. Few studies have attempted to design more ecologically valid experiments that might better reflect the usual context in which voluntary actions are made.

This thesis will attempt to address these issues, using established and novel paradigms to investigate two aspects of voluntary action: volition and agency, in health and disease.

1.3. Introducing volition

1.3.1 Volition and free will.

Volition, or the intention to act, is an intuitively familiar experience, yet its specifics remain a source of debate among philosophers, artists and scientists (Roskies, 2010). Written records show that some early Greek philosophers, notably early Materialists such as Epicurus, attempted to describe the worldly origins of human action without recourse to 'divine influence' (Bobzien, 2000).

The concept of 'free will' is not directly addressed in this thesis. All experiments presented here presuppose that volition arises as part of a deterministic process of causes and effects within the brain.

1.3.2. What volition isn't.

It is unusual to speak about volition for reflexes and automatic bodily functions. Breathing, blinking, and sneezing are common examples of actions that seem to happen without any effortful or conscious intervention. Even movements that are typically deliberate, such as raising the leg or arm, may be induced by percussion or other kinds of external stimulation, such as an imperative 'go' signal in an experimental context. The less a movement can be derived from a single, immediate sensory cause, the more likely it seems to be described as voluntary. This aspect of volition has neatly been referred to as a "freedom from immediacy" (Shadlen & Gold, 2004).

Following from this tradition, volition in the current work is taken to be an endogenous intention to perform some action. Endogenous in this case meaning that it arises

without the direct and temporally contiguous influence of an identifiable external stimulus.

A further important distinction must be made between volition, as described above, and 'agency' which will be introduced separately. Experiments investigating agency will consider the entire subjective experience of producing an action, including execution of the movement and monitoring its subsequent effects. Volition here will be restricted to the intention to act, leading up to the moment of action execution.

1.3.3. Studying volition – practical considerations

The study of volition through observations of behaviour presents a critical design challenge. To provide useful data, participants must respond in some way to the experimental environment, usually by producing an action. In many experimental paradigms, an explicit instruction or "go-cue" is provided. However this method of prompting an immediate action is not compatible with volition as described here. Actions must be elicited without compromising the endogenous character of the intention to act. Participants must feel that they are free to act whenever they choose. Immediate, punctate stimulus cues are thus considered unsuitable.

One solution to the problem of eliciting actions without a specific go-signal is through the use of **partly-directed actions**. In such paradigms, the participant is given explicit instructions to produce an action *at a time of their choosing* (for example Libet et al., 1983). In this manner, participants can still produce many trials of equivalent movements in response to non-signal stimuli and their responses may be recorded.

However, in some experimental designs it remains preferable to cue an action at a precise moment, e.g. when the length of individual trials should be kept consistent. In

such cases, participants may be instructed to choose *what* they do instead of *when they do it* (e.g. Haggard & Eimer, 1999). In this method, actions are cued but preserve a critical aspect of volition: the need for the participant to make an action-related decision.

In some designs, participants may decide to make no action at all. **Inhibiting movement** is a special case of partly-directed action. Inhibition is a particularly difficult cognitive phenomenon to study empirically due to the absence of any measurable output. Eliciting the absence of voluntary action may be achieved however, by instructing the participant to withhold an action that would otherwise normally occur as part of a sequence (e.g. Misirlisoy & Haggard, 2014).

The use of partly directed actions as a means to study volition is a convenient compromise, but not without its limitations. Paradigms in which participants choose *when* to act are typically bound by a maximum time period. This time period is typically very long and not usually exceeded in practice. Nevertheless, the limit curtails the freedom to act completely at will. Similarly, the very existence of instructions within an experiment diminishes volition, particularly in repetitive tasks, such as those frequently employed in laboratory paradigms. These criticisms must certainly be acknowledged and considered in good experimental design. Nevertheless, partly-directed actions remains a useful, if imperfect, tool for describing behavioural and neural features that accompany simple voluntary movements.

1.3.4. Simple models of volition

A simple model of volition may be drawn intuitively from everyday experience. In this model, there is an intention that temporally precedes an action (figure 1-1).

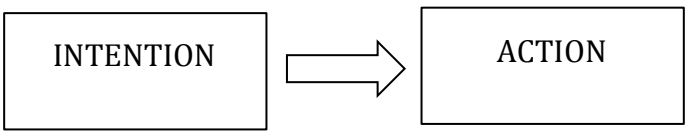

Figure 1-1: A simple model of volition

This simple model does not account for any function of the brain, much as dedicated models of musculoskeletal movement typically do not consider intentions. Evidence-based models of volition that included both intention and brain activity were not considered for many years. This was partly due to the inability to precisely record brain activity at the very precise time scales required.

The pervasive feeling from everyday experience is that the intention (whatever its antecedents) is the first event to occur, so brain activity should likely arise after, possibly even as a result of, intention (figure 1-2). This is the naïve dualism of everyday life as a conscious agent, and is deeply embedded in our culture.

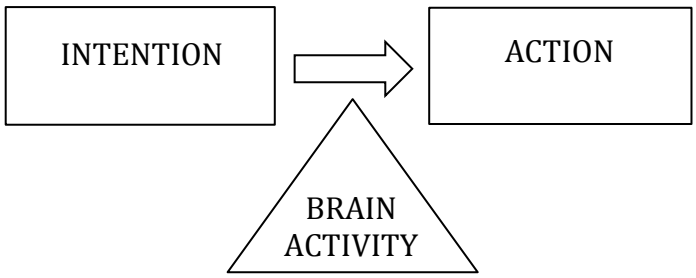

Figure 1-2: An intuitive model of volition in which brain activity is assumed to occur between intention and action

1.3.5. The Libet experiment

In 1983, Benjamin Libet and colleagues reported the results of an experiment that provided electrophysiological evidence against the naïve dualist model. In their experiment, participants made a simple action (upwards flexion of the wrist) while

watching a purpose-designed clock with a single rotating indicator. This method of using a chronometric device to report punctate events dates back to the experimental designs of Wundt (1907). Participants were instructed to flex their wrist at a time of their choosing and, depending on condition, to note the position of the clock when:

- They first felt the urge to move (referred to as a 'W' judgement)
- They actually flexed their wrist ('M' judgement)

In a separate condition, participants did not perform an action but receive a sensory stimulus (tactile contact to the hand) and reported the time they perceived this to occur ('S' judgement).

At the same time that this simple behavioural task was being performed, the authors recorded EEG measurements to determine the onset and nature of the so-called 'readiness potential' (RP). Also known as the '*Bereitschaftspotential'* (Kornhuber & Deeke, 1965) the RP is a reliable, consistently negative-going component that precedes actions. The naïve dualist model predicts that the RP will not begin until after the participant-reported W judgement, i.e. *first* the intention is formed *and then* the brain begins to prepare at an appropriate interval before movement. Libet and colleagues found evidence contrary to this prediction.

The results obtained by Libet et al. (1983) showed that the onset of the RP measured at the vertex, where it was maximal, occurred significantly in advance of the intention to act as reported by the participant (W judgement). RP onset was typically observed more than 1 s before the actual movement. By contrast, participants' W judgements were on average only about 200ms before the action occurred (see figure 1-3 below). This was taken as clear evidence that brain activity preceded intentions and not vice versa.

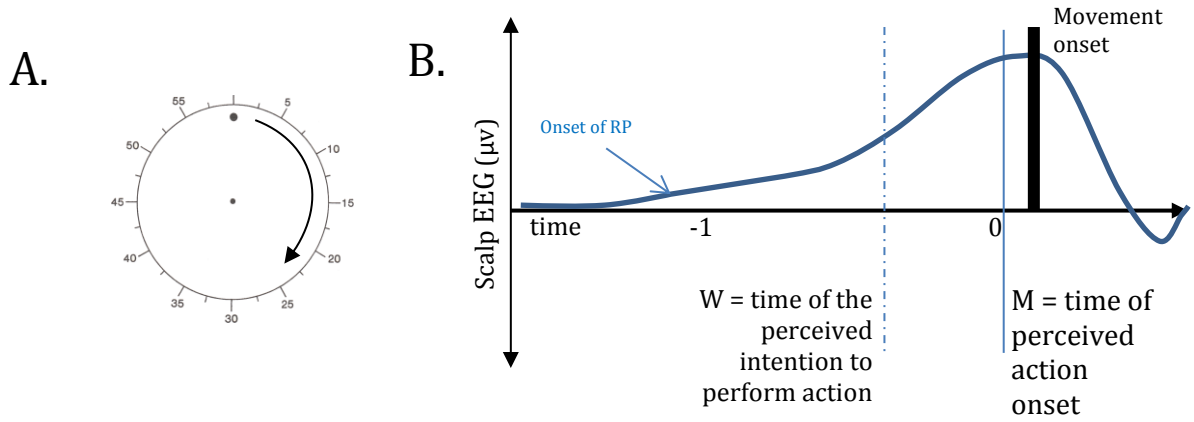

Figure 1-3: The Libet paradigm. A: an example of a clock stimulus used by participant to note the time of intention and action events. B: a schematic RP trace showing typical scalp voltage readings over time, with critical experimental events (W and M judgements) marked.

The method described by Libet is not without its shortcomings (see Banks & Pockett, 2007 for a review). It is unclear how participants deploy and divide their attention between the clock face, the moving indicator, and other elements such as the button press. Furthermore, participants' reports may also be unreliable due to retrospective reconstruction and individual strategies for judgement. While these criticisms are valid, the paradigm has enjoyed more than three decades of replication, with results that are seemingly robust to a variety of contexts, materials and populations (Eimer & Hagard, 1999; Moretto et al., 2011; Ganos et al., 2014). Results from experiments that employ this method must therefore be interpreted with caution, but represent an important means of insight into the underlying mechanisms of voluntary action. Results in Libet-type experiments suggest that participants can report a subjective experience related to a voluntary action, but do not in themselves show what processes give rise to this experience.

1.3.6. Pre-movement cortical activity – the Readiness Potential

The precise definition and nomenclature of pre-movement cortical potentials has historically varied considerably by author and by research question (Kornhuber & Deecke, 1965; Shibasaki et al., 1980; Tarkka & Hallett, 1991; Cui & Deecke, 1999). Relevant to the aims and focus of this thesis, investigations on self-initiated actions generally refer to a readiness potential (RP) or *bereitschaftspotential*. The two main features of the RP are generally considered to be: an early, negative-going segment with a shallow slope that begins around 2s prior to movement onset, and a later segment with a steeper slope beginning at about 400ms prior to movement onset (Shibasaki & Hallett, 2006).

The neuroanatomy of the RP has been elucidated by various methods including principal component analysis (PCA) and EEG dipole source analysis (Toma, 2002), as well as recording from electrodes implanted in patients being treated for intractable epilepsy (Ikeda & Shibasaki, 2003). The consensus from such investigations is that the RP is first detectable in the anterior part of the supplementary motor area (SMA) known as the pre-SMA, then in the SMA proper, followed by the premotor cortex. The late segment of the RP is centred on the primary motor area (M1) (Neshige et al., 1988), which is unsurprising given its immediate temporal proximity to movement onset.

Interestingly, the two main components of the RP may each be influenced by different factors. Learning, preparation and speed of movement have been shown to influence the size (learning) and onset (preparation, speed) of the early segment (Lang, 2003). The later component may be influenced by precision (Simonetta et al., 1991) and complexity (Kitamura et al., 1993). Both may be affected by pathology (see below section 1.3.9.).

A critical consideration when extracting RP traces from EEG data is determining a suitable baseline. The baseline should precede both intention and RP onset, but no independent evidence is available to determine when an RP begins. A common practice is to select a short window from the beginning of recordings and take the average amplitude in that window as a baseline. However this makes a strong assumption that activity at recording sites is somehow at a consistent default or rest state and that the RP has not yet begun. This assumption and a potential correction are addressed in Chapter 3 of this thesis.

1.3.7. Wider neuroanatomy of voluntary actions

The readiness potential does not occur in isolation, nor is it the only important input to M1. Research on monkeys shows that the pre-SMA and SMA proper receive inputs from both the basal ganglia and, to a much lesser degree, the cerebellum (Akkal, Dum, & Strick, 2007). In humans, Loukas and Brown (2004) used recordings from implanted electrodes in the subthalamic nucleus (STN) to successfully predict self-paced movements. Figure 1-4 below shows putative motor preparation pathways. The implications of a subcortical-cortical loop in patients with reduced neural output from the basal ganglia (e.g. in Parkinson's disease) is discussed below in section 1.3.9.

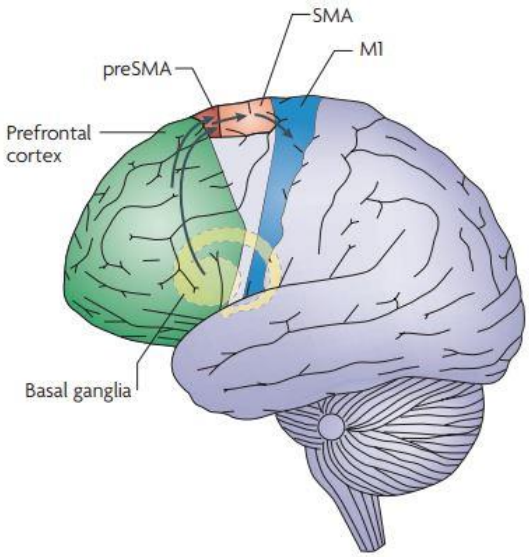

Figure 1-4: A pathway of cortical activity preceding onset of a voluntary action, showing projection from the basal ganglia and prefrontal cortex to the SMA (initially the pre-SMA) and then to the primary motor area (M1). Adapted from Haggard (2008).

1.3.8. From brain activity to a subjective experience of volition

In an attempt to more clearly describe how brain activity may give rise to a subjective feeling of volition, an influential model was proposed by Hallett (2007). Hallett's model is largely informed by primate experiments in which movements are triggered when the activity of relevant cell networks reaches a specific threshold. This has been observed in the frontal eye fields for saccadic movement (Schall, 2002) and in the motor cortex for limb movements (Lecas, Requin, Anger, & Vitton, 1986). According to Hallett, volition arises when motor preparation reaches a specific initial threshold and the action itself arises when motor preparation reaches a later threshold. The later action threshold represents a 'point of no return', after which the movement necessarily occurs.

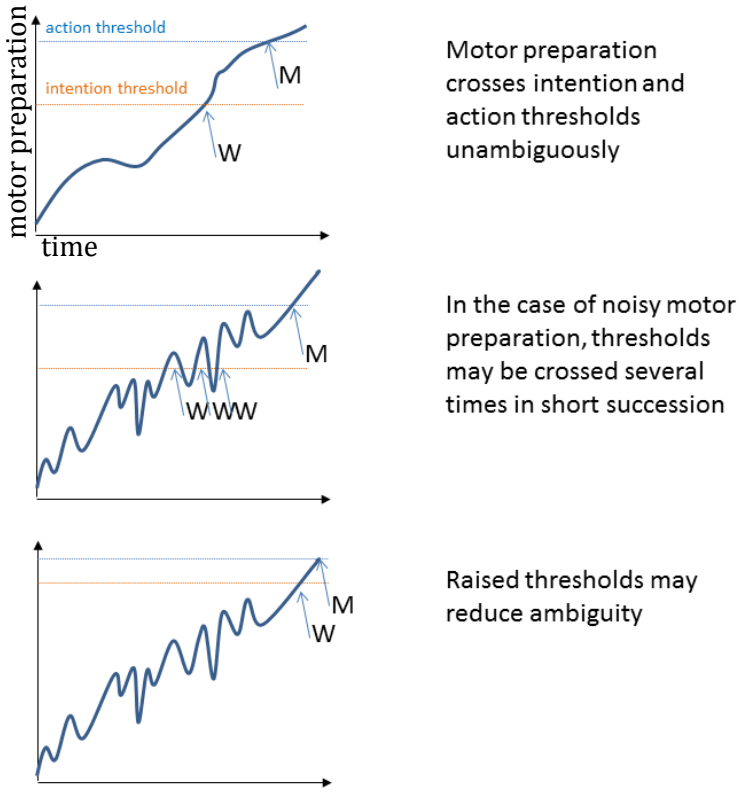

Figure 1-5: Modelling awareness of intentions and actions as signal detection. Y-axes represent accumulation of neural activity in premotor areas (motor preparation) preceding a specific action. X-axes represent time in the moments before the action takes place. Where initial motor preparation is noisy (lower two frames) raised threshold may provide a less ambiguous signal. Here W represents awareness of intention and M represents awareness of the movement 'trigger' or point-of-no-return.

Importantly, the accumulation of motor preparation is thought to occur in frontal motor networks (Hallett, 2007; Fried et al., 2011), as part of the subcortical-cortical loops described earlier. Given that these are the specific areas that precede action initiation from M1, the timings of W and M judgements in the Libet paradigm map easily onto Hallett's model. The first threshold, at which one becomes aware of the intention to move, is the moment reported as the W judgement. The second threshold, at which one becomes aware of movement onset, is reported as the M judgement.

Results in a Libet experiment may therefore provide insight into the underlying features of motor preparation preceding a voluntary action. W and M judgements that are temporally close might indicate rapidly accumulating motor preparation, which would

appear as a steep curve when plotted as in figure 1-5. Similarly, W and M judgements that are temporally distant may indicate a shallow curve, representing a slow accumulation of motor preparation. However, changes in the perception of W and M events may be due to other factors affecting motor preparation, such as neural noise.

In this model, W and M are assumed to represent different thresholds on the same underlying signal, thus affected by the same neural noise. The presence of noise in the motor system results in periods of both increasing *and decreasing* motor preparation. Neural activity may therefore cross the threshold for volition (W) several times and by several networks (see figure 1-5, middle plot) before reaching the action threshold (M). It therefore follows that the W judgement arises in a more ambiguous neural context than the M judgement and is thus more variable. This is indeed typically observed in Libet-type experiments, including those conducted within the present work (see chapters 2 & 3). One may also consider that thresholds are not fixed and may be raised or lowered contingent on the neural context. Critically, this model predicts altered thresholds in patient groups where motor-related neural noise is pathologically disturbed.

1.3.9. Disordered volition

The subjective nature of volition and the absence of any clear criteria to define precisely when it occurs make speaking about 'disordered' volition difficult. However, certain neurological pathologies are characterised by disorders of voluntary movement. The neural abnormalities that accompany such pathologies are by now well known. One might reasonably suppose that divergences in these underlying mechanisms should manifest as altered experiences of volition. Patients who show abnormal patterns of

action initiation, maintenance and termination are therefore a useful population in which to study the phenomena associated with volition.

Parkinson's disease (PD) is characterised by reduced and slower voluntary movements compared to healthy populations (Goetz et al., 2008). Cognitive impairments in PD include an increase in impulse control disorders (Weintraub, David, Evans, Grant, & Stacy, 2015), apathy (Starkstein et al., 1992) and depression (Cummings, 1992).

The primary site of pathology is the substantia nigra, where PD patients show severe loss of dopaminergic projection neurons relative to healthy individuals (Braak & Braak, 2000). However, even from the earliest stages of the disease, this loss is associated with impaired function beyond the substantia nigra. The systems affected are varied. Research has largely focussed on disorders in dopaminergic pathways in the limbic system, however other cell types including cholinergic, GABAergic and glutaminergic neurons are also typically implicated (Jellinger, 1991).

Lesions in the limbic system, particularly the transentorhinal region, have wide-reaching consequences in PD (Braak et al., 1994). Damage in entorhinal and hippocampal areas, as well as the amygdala, interrupts normal communication from these sites to cortical and subcortical regions. Learning, memory, emotion regulation and motor function are all typically affected (Braak et al., 1996).

Relevant to the experience of voluntary actions, diminished dopaminergic projection to the SMA have been shown to reduce both the occurrence and speed of self-generated movements (Jahanshahi et el., 1995). Given the involvement of the SMA and pre-SMA, it

seems reasonable to assume that neural markers associated with the motor symptoms of PD may be detectable in changes to the readiness potential (RP).

RP studies of PD have typically measured the amplitude or latency of the negative going pre-motor potential. Results have differed considerably between studies (Barrett, Shibasaki, & Neshige, 1986; Dick et al., 1987; Jahanshashi et al., 1995; Shibasaki et al., 1980), however this may reflect methodological differences, particularly in the administration of dopaminergic medication. A recent review (Georgiev, Lange, Seer, Kopp, & Jahanshahi, 2016) suggests that the most consistent finding from among 14 relevant studies was a reduction in the amplitude of the early stage RP. The authors further noted that the diminished amplitude could be increased by the administration of dopaminergic medication (Feve, Bathien, & Rondot, 1992; Dick et al., 1987).

A recent study (Tabu et al., 2015) examined the awareness of action intentions in PD using a Libet task. Participants were elderly (mean age of 67 for males, 70 for females), mildly impaired according to established criteria (Goetz et al., 2008) and completed the task having not taken dopaminergic medication for at least 12 hours. Patients showed a much later W judgement (awareness of intention to act) relative to healthy controls, while M judgements (awareness of acting) and S judgements (awareness of a sensory event) were unaffected. The authors concluded that due to the relatively mild impairment of the cohort, these results were unlikely to be the sole product of cognitive deficits.

It is tempting to conclude that reduced levels of dopamine (DA) may have driven delayed intention awareness alongside diminished movement control. Tabu and colleagues did not test patients while on dopaminergic intervention, so results from their study cannot be used to defend this viewpoint. Furthermore, this hypothesis is not

supported by results from patients with Tourette's syndrome (Moretto et al., 2011), described below.

Tourette's syndrome is characterised by involuntary movements and vocalisations called tics that typically appear in late childhood and often continue into adulthood. The ability to control tics varies both *between* individuals and *within* individuals over the lifespan (Robertson et al., 2009). Tics are a particularly unique case of 'unwanted' movement in that they are not described as fully automatic (like reflex movements). Patients often report 'urges' prior to a tic, and many patients learn to control tics as they move from adolescence into adulthood (Jackson et al., 2011). Tics are therefore experienced as involuntary, but are mediated by the voluntary motor system. It is particularly interesting to note that RPs may be observed prior to tic movements, however this is rare and the RPs are shorter in duration compared those of fully intended movements in healthy controls (van der Salm et al., 2012).

Ganos and colleagues (2015) predicted that increased neural noise preceding movements in Tourette's syndrome would result in a later subjective experience of action initiation. They reasoned that a less reliable signal around the usual time of RP onset would require a more conservative threshold for the detection of volition. Indeed, Moretto and colleagues (2011) previously observed significantly later action awareness in Tourette's syndrome compared to healthy controls. Ganos et al., using the Libet method in a study on 27 adolescents, failed to replicate Moretto et al., but found that patient reports of volition were indeed influenced by features of the pathology. Specifically, although age and severity of tics was not systematically related to a delayed awareness of intention, patients who experience strong pre-tic urges also reported a later intention to act. Similarly, patients who had good control of their tics, showed a

much earlier awareness of intention. The authors concluded that volition arises as the brain learns to discriminate neural noise related to pre-motor preparation.

Importantly, Tourette's syndrome is a *hyper*-dopaminergic condition; so reduced DA levels cannot explain delayed awareness of intentions.

It therefore remains to be determined whether disorders implicating dopamine, such as Tourette's syndrome and PD, are associated with delayed intention awareness due to one common or several distinct neural processes. One of the objectives of this thesis is to investigate this question by comparing the action awareness of PD patients with healthy controls. Critically, the presence of DA-related symptoms in patients will also be manipulated in three separate experiments using medication and deep brain stimulation (DBS). These interventions will be more fully introduced in the relevant sections of chapters 2 & 3.

1.3.10. Volition – conclusions

Before introducing the *Sense of Agency,* a related but distinct aspect of voluntary movements, it is useful to consider volition in light of the work reviewed thus far.

Volition is a key subjective experience in 'making things happen' and separates considered, planned actions from reflexes and other automatic movements. The brain areas associated with volition, and the patterns of activation preceding a voluntary movement are by now quite well described. However, it remains to be seen how the neural processes that accompany volition are involved in abnormal expressions of action awareness in certain pathologies, including Parkinson's disease. It is hoped that a better understanding of disordered volition will lead to a greater understanding of volition generally.

Chapters 2 and 3 of this thesis will describe experimental work investigating volition. Brief descriptions of the research aims of these chapters can be found in section 1.5. of this introduction.

1.4. Introducing the Sense of Agency

1.4.1. Definitions

Volition describes an intention to act, prior to physical movement. Yet the subjective experience of 'doing something' extends beyond the moment of action onset. The everyday feeling of controlling one's actions is not just about willing them to occur. When our actions produce effects in the environment, we feel that we cause those too. This cluster of feelings is often referred to as the Sense of Agency (SoA hereafter) (Gallagher, 2000; Haggard & Chambon, 2012). SoA is the interplay of volition, action and outcome, and a ubiquitous attribute of everyday life. Linking one's actions to their effects informs learning and adaptation, and has been critical in the development of man-made civilization. Long-term planning and organisation does not seem possible without the ability to appreciate the temporally distal consequences of one's actions. Agriculture, settlement, exploration and exchange have all relied on a developed SoA.

Synofzik, Vosgerau, and Newen (2008) describe two main aspects of agency: the **feeling of agency** and the **judgement of agency**. The latter refers to an explicit declaration that an effect was or was not caused by one's actions. The *feeling* of agency, on the other hand, is described as a basic, pre-reflective sensorimotor experience, which exists independent of explicit judgement. These two phenomena cannot be measured in the same way, and each presents distinct challenges in experimental design.

1.4.2. Measuring agency

Pre- and post-reflexive measures of agency may not be measuring precisely the same thing. Dewey & Knoblich (2014) recorded both implicit and explicit measures in a

simple agency task. They found that explicit and implicit measures did not significantly correlate. Therefore caution is required when interpreting and comparing results both within and across experiments.

Explicit judgements of agency may be measured experimentally by asking participants to judge whether or not they caused an effect (e.g. Spengler, Von Cramon, & Brass, 2009; Maeda et al., 2012). Alternatively, participants may indicate the degree to which they feel they caused the effect (e.g. Wenke, Fleming, & Haggard, 2010; Sidarus, Chambon, & Haggard, 2013). A critical consideration in experiments such as these is the presentation of an alternative 'cause' that may realistically be considered by participants. Depending on research questions and practical considerations, this may be presented as a fellow participant, a confederate or an automated entity such as a computer programme or computer-controlled avatar. However, explicit judgements of control or agency are influenced both by performance bias (Metcalfe & Greene, 2007), and by a general self-serving bias (Bandura, 1989). A confounding effect of errors on explicit agency judgements therefore seems inevitable.

To avoid these potential confounds, and for clarity of interpretation, the experimental work in this thesis will be restricted in focus to the pre-reflexive *feeling* of agency.

The feeling of agency is not directly measurable. Quantifying a subjective feeling of control normally requires finding a robust, replicable proxy measure. Furthermore, the chosen measure should be generalizable. That is, SoA should be reasonably independent of the effector used for control (e.g. hand, eye movement, speech), and reasonably independent of the modality of the outcome (sound, light, emotion).

Chronometric approaches are frequently employed. Participants use an external timekeeping device (typically a purpose-designed clock) as an objective reference point for stating when they made an action. This method is based on the studies of volition described above (specifically the 'Libet method', see section 1.3.5.). Critically however, agency research also makes use of 'outcome events' that follow predictably and reliably from the action. These are typically simple auditory or visual stimuli. The perceived time of a sensory outcome is often measured alongside the perceived time of the action.

The use of chronometry is a convenient method of obtaining implicit measures of agency. However, alongside independence of effector and modality, SoA is also not exclusive to a specific timescale. For example, deciding to press either the left or right key in an experimental task takes less than a second, whereas the decision to buy a house takes place over a much longer period. Ideally, the method used to measure agency should apply across several timescales. At present, the methods used in agency research may not adequately fulfil this criterion as they typically work in sub-second timescales. Notable exceptions have been provided by Engbert, Wohlschläger, Thomas, & Haggard (2007) and by Humphreys and Buehner (2009), who employed time *interval* estimation in place of clock-based judgements. Their results indicated that reliable binding effects could be observed at super-second intervals (up to 4 s).

Restricted generalisability of timescale is accepted as a limitation of the experimental work in this thesis, which makes extensive use of a chronometric method known as Intentional Binding.

1.4.3. The Intentional Binding paradigm

As described, earlier, paradigms that make use of the Libet method may suffer from suboptimal reliability of participant responses (Banks and Pockett, 2007; and see section 1.3.5.). Contemporary experiments investigating SoA are typically designed to safeguard against such criticisms. This is usually achieved by measuring the shift in judgements from one condition or group to another. In this case, inaccurate judgements are not problematic provided they are consistently inaccurate.

The "Intentional Binding" paradigm (Haggard, Clark, & Kalogeras, 2002; Walsh & Haggard, 2013; for a review, see Moore & Obhi, 2012) has been proposed as a robust proxy measure for sense of agency. Participants report the time of an action (e.g. a key press) or of an outcome (e.g. an auditory tone). Taken in isolation, these judgements serve as a baseline. In operant conditions, the outcome follows the action rapidly and reliably. In these conditions, the perceived time of the action tends to shift towards the outcome. Similarly, the perceived time of the outcome tends to shift towards the action (figure 1-6.). Critically, this effect is not observed for involuntary movements (Haggard, Clark, & Kalogeras, 2002). Intentional binding is a convenient, implicit, proxy measure of agency. It also may be related to perception of causation more generally, though this remains controversial (Buehner & Humphreys, 2009; Buehner, 2012; but see Cravo, Craessens, & Baldo, 2009; Cravo, Craessens, & Baldo, 2011). There may be many possible causes of temporal attraction, however experimental designs that compare results between critical conditions can isolate the influence of intention.

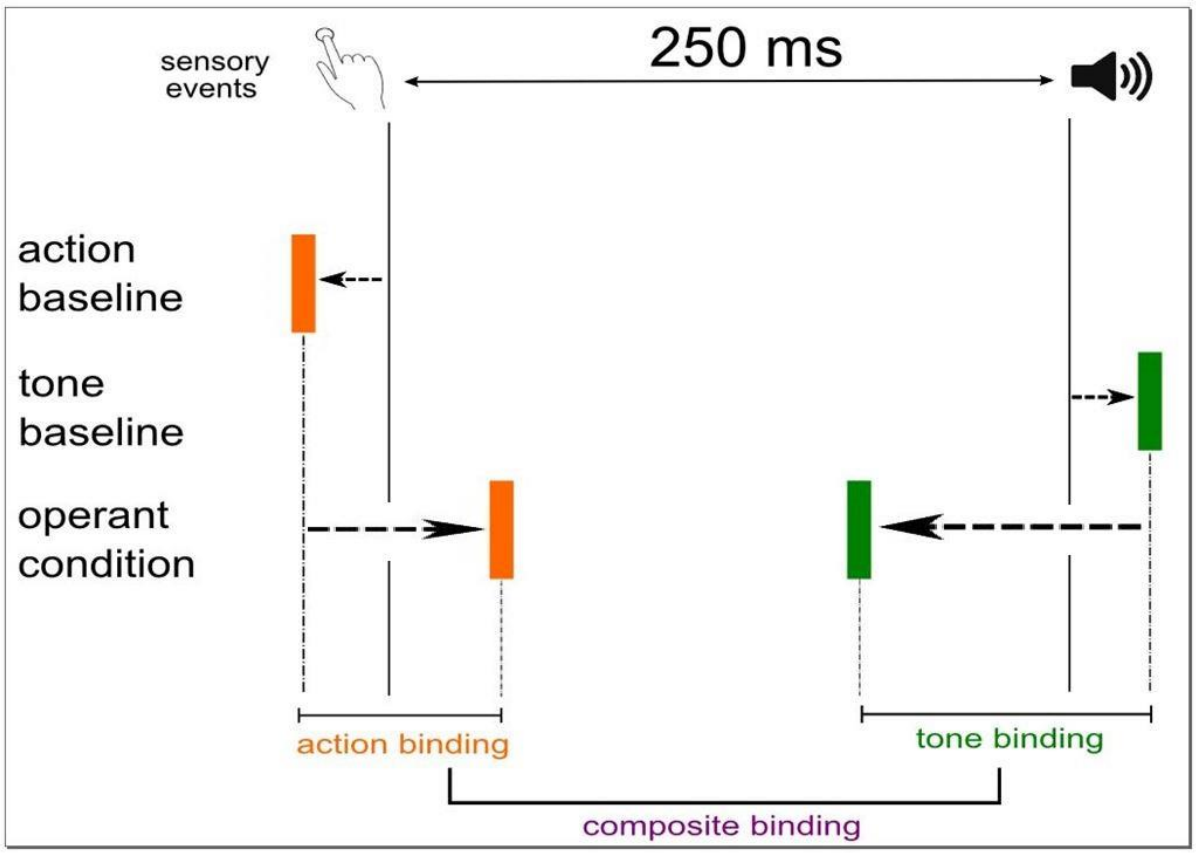

Figure 1-6: The intentional binding paradigm. Judgements in baseline conditions for actions occurring alone or tones occurring alone may not coincide with the actual time of the corresponding event. When the action causes the tone after a delay, the perceived times of action and tone are shifted towards each other. Reproduced from Di Costa, Théro, Chambon, & Haggard (2017).

1.4.4. Interpretation of binding

An important research question, hitherto unresolved in the literature, is precisely what drives binding. Research suggests a role for various candidate mechanisms.

Cue integration models suggest that information from multiple sources is combined in a weighted average. Ernst and Banks (2002), for example, asked participants to estimate the height of a stimulus given visual and haptic information. Reliability of information was manipulated by increasing variance. Results showed that the influence of each modality on the final estimate depended on the reliability of that information.

Applying this model to Intentional Binding implies that action and outcome serve as cues. When one cue is less reliable or salient, it may be 'captured' by the other stimulus. Capture in this sense would be observable by increased perceptual shift. Wolpe et al. (2013) provided partial evidence for this interpretation. The authors manipulated the reliability of auditory outcomes by playing sounds at volumes near or far from participants' perceptual thresholds. They found that, as predicted by cue integration, action binding was reduced when volume was near the perceptual threshold. In other words, the action event was not captured as strongly by the auditory outcome when this was weakly perceived. Tone binding appeared unaffected by the manipulation, suggesting that the mechanisms driving action and tone binding are at least partially distinct.

Association formation has been proposed as an alternative candidate mechanism underlying binding. Evidence for this hypothesis was reported in a binding experiment that manipulated action-outcome contingency (Moore, Lagnado, Deal, & Haggard, 2009). Participants chose whether or not to press a key on each trial. For some participants, the probability of a tone occurring after a keypress was set at chance (50%). For other participants, the probability was set at 75%. Contingency was defined in this experiment as the probability of an outcome given an action, minus the probability of the same outcome given no action. Therefore, the authors separately manipulated the probability that a tone would occur in the *absence* of a keypress. Participants in each condition therefore experienced both *contingent* and *non-contingent* blocks of trials.

Results indicated that binding was indeed subject to contingency. In the low probability group, contingency was associated with stronger binding only on trials where the tone did in fact occur. Conversely, in the high probability group, contingency was associated

with stronger binding in trials where the tone didn't occur. These results add weight to the theory that association formation is implicated in binding, but also suggest that binding may involve both prospective and retrospective components.

1.4.5. Interpretation of binding - prospective and retrospective mechanisms

Research suggests that the mechanisms underpinning intentional binding may have distinct directions of effect. Wegner and Wheatley (1999) discuss how participants can, in certain circumstances, assume responsibility for outcomes they did not actually cause. This suggests that agency has a retrospective component, independent of any *a priori* intention.

A prospective component of agency has also been identified. Moore and Haggard (2008) varied the predictability of outcomes in learned action-outcome sequences. They found that when auditory outcomes followed actions infrequently, binding was detected only on trials where the tone occurred. When tones followed actions frequently, binding was also observed when tones were absent. The formation of reliable associations therefore appears to be a critical component in binding.

A similar study investigated this effect in both schizophrenic and healthy subjects (Voss et al., 2010). In a high probability condition, tones followed key presses 75% of the time. In a low probability condition, tones followed key presses only 50% of the time. The prospective aspect of binding was replicated in the control group, where stronger binding was observed in the high probability trials, even when the tone was absent. The patient group did not show this effect. Interestingly, a strong retrospective influence was detected in patients, who showed stronger binding on trials where the tone occurred in the low probability condition. This effect was not observed in the control

group. This pattern of effects suggests that when a reliable predictive mechanism is lacking, retrospection may make take on greater importance.

It seems clear that both predictive and retrospective mechanisms can influence the occurrence and strength of intentional binding. It is unclear whether these mechanisms interact or remain exclusive from each other. More research is also required to adequately determine the factors that contribute to prospective and retrospective influences on implicit measures of the sense of agency.

1.4.6. Desired and undesired outcomes

Our actions do not always produce the effects we would like. In an ideal world, we would only cause positive, desired changes in our environment. Occasionally however, negative, undesirable effects occur. Only a few studies have attempted to link the sense of agency with outcome valence. Takahata et al. (2012) associated auditory tones with monetary rewards or penalties. Their results showed reduced binding for trials that resulted in a penalty compared to neutral or rewarded trials. Yoshie and Haggard (2013) used valenced human voices as outcomes and found that binding was significantly reduced when actions produced negative outcomes. Interestingly, in both cases only negative outcomes produced an effect. Positive and neutral outcomes were not differentiated in binding measures. Furthermore, only composite binding measures (in which action and tone binding are added) were reported in both cases.

It therefore remains to be determined how valenced outcomes may affect action or tone binding separately. Manipulation of desired and undesired outcomes is a hallmark of decision-making research. This thesis will therefore make use of established decision-making paradigms, introduced below, to investigate the sense of agency.

1.4.7. Reward-based decision-making

Studies investigating reward-based choice behaviour typically involve explicit goals. Goals may include accumulating points, earning money, or simply causing a pre-defined stimulus to manifest. Importantly, these studies are not usually concerned with any specific motor movement. Rather, the focus is on the explicit choice being made. A critical question is why one alternative would ever be chosen over another.

1.4.8. Reinforcement learning

Reinforcement learning offers a parsimonious explanation for choice behaviour. The basis of reinforcement learning is operant conditioning. An action that precedes a favourable outcome is likely to be repeated (Skinner, 1938). A critical component of this mechanism appears to be temporal proximity. Different kinds of rewards have successfully induced reinforcement learning in humans. Examples that have withstood decades of replication include monetary and food rewards (see Sescousse, Caldú, Segura, & Dreher, 2013 for a review).

The neural basis for reward processing is well-established. Old and Milner (1954) directly stimulated the brains of rats whenever they pressed a lever. Stimulation in the septal area was enough to drive perseverative pressing of the lever. These results were striking because the rodents did not experience the reward stimulus with their senses, only as neural stimulation. In humans, Knowlton, Mangels and Squire (1996) found that PD patients with neostriatal damage were unable to learn reward contingencies. This occurred despite patients reporting accurate memory for the training events. These results were compared with amnesiac patients who retained no memory of training, yet learned the reward contingencies.

Dopaminergic midbrain neurons appear to play a key role in the so-called 'reward circuit' (Haber & Knutson, 2010). Drevets et al. (2001) used Positron Emission Tomography (PET) to measure drug-induced DA release in humans. DA increase in the ventral striatum positively correlated with subjects' reported 'euphoria'. These results were compatible with previous animal research.

Sometimes it is not a reward itself but the expectation of reward that matters. Reinforcement learning relies on the fact that subjects associate sensory events with outcomes. When the predicted consequence of an event fails to arise, a negative prediction error occurs. Conversely, a *positive* prediction error accompanies the delivery of an unexpected reward.

Prediction errors are often studied using probabilistic reward schedules. McClure, Berns, & Montague (2003) trained human subjects to expect a juice reward following presentation of a light. On some trials, juice was withheld following the light. On other trials, juice was delivered without the light. Blood-oxygen level dependent (BOLD) changes were measured with fMRI in the period before subjects would normally expect delivery of a reward. Negative prediction errors (expected juice, no juice delivered) were accompanied by reduced activity in the putamen. Positive prediction errors (no expectation of juice, juice delivered) were accompanied by increased activity in the putamen. Both contrasts were made to 'normal' events (juice expected and delivered, or no expectation of juice and no juice delivered). Single unit recordings of dopamine neurons in the macaque have shown a similar pattern of results (Hollerman & Schultz, 1998; Bayer & Glimcher, 2005).

Prediction errors are clearly important for learning in dynamic environments. Inducing prediction errors in decision-making experiments is often achieved using reversal learning tasks.

1.4.9. Reversal learning

Reversal learning paradigms have been used extensively in psychological research. The paradigms typically require the participant to track occasional changes in action-outcome mappings, and adjust action selection accordingly (Rolls, 2000; Cools, Clark, Owen, & Robbins, 2002). The critical feature of such tasks is a transition, either gradual or abrupt, in the value of choices. Participants must therefore constantly monitor outcomes to ensure optimal decision-making.

Reversal learning is normally distinguished from tasks where probability schedules shift constantly and independently. The two-stage Markov task is one such example (Daw, Gershman, Seymour, Dayan, & Dolan, 2011; Wunderlich, Smittenaar, & Dolan, 2012; Smittenaar, Prichard, FitzGerald, Diedrichsen, & Dolan, 2014). In these tasks, the chance of reward shifts by a small amount on each trial, but an increase in probability is as likely as a decrease. In other words, probabilities are shifting, but not necessarily reversing.

A study by Cools et al. (2002) offers an illustrative example of the reversal-learning paradigm. In this task, participants were presented with two abstract patterns on a computer screen. In the acquisition phase, choosing one of the patterns usually led to a positive outcome (green smiley face). Choosing the other usually led to a negative outcome (red sad face). During the reversal phase, the probability schedules reversed after between 10-15 'optimal' choices. The authors identified two kinds of error in this

task: *reversal errors* were sup-optimal choices made immediately following a reversal, i.e. continued reliance on the previously rewarded choice. *Probability errors* were optimal choices that were followed by negative feedback. Critically, the subject cannot distinguish whether any individual error trial is a reversal or probability error.

Research shows that patients with deficits in planning behaviours are typically impaired in performance on reversal-learning tasks. Ersche et al. (2008) found that stimulant-dependent individuals (cocaine or amphetamines) took significantly longer to reach a reversal criterion than healthy controls. Patzelt, Kurth-Nelson, Lim & MacDonald (2014) showed that this is likely to be due to excessive switching, at least in cocaine users.

Reversal-learning paradigms are ideally suited to examine goal-directed decision-making for two reasons: firstly, the ability to separate action-outcome associations from outcome valence, and secondly, the sensitivity of the task to planning and strategizing deficits and the ability to categorize changes in learning rate.

It remains to be determined how performance on such tasks interacts with measures of SoA. Indeed, the fields of motor control and goal-based decision-making have evolved separately in the literature. Chapters 4, 5, 6 and 7 of this thesis will describe work that represents a first attempt to combine these distinct research traditions.

1.4.8. Neural correlates of the Sense of Agency

Depending on individual definitions and methods, SoA may implicate a number of cognitive processes including volition, action monitoring and time perception generally. Delineating specific neural substrates related to agency is therefore not trivial.

Sperduti and colleagues (2011) conducted a meta-analysis including more than 200 subjects from 15 neuroimaging (PET and fMRI) studies of agency. In these studies, participants provided judgements or implicit attributions of agency to the self or to an external source (computer programme, avatar etc). Comparing across these two conditions, the authors found a dissociable pattern of neural behaviour. Non-agency was associated with activity in the inferior parietal lobe and superior temporal gyrus, as well as the dorsomedial prefrontal cortex. Self-agency was most strongly associated with activity in the insula. These results are perhaps unsurprising, as connectivity in related areas has also been described in neuroimaging studies of volition (see section 1.3.6.). 'Intention information' would be a necessary component of models of agency based on the comparison of desired and actual outcomes, such as the comparator model (Blakemore, Wolpert, & Frith, 2002).

Evidence from non-invasive brain stimulation experiments has also shed light on agency in the brain. Khalighinejad, Di Costa and Haggard (2016) conducted a meta-analysis of seven experiments investigating the effect of anodal transcranial direct current stimulation (tDCS) of the prefrontal cortex on intentional binding. Results suggested that stimulation increased the temporal binding of actions to their outcomes. Critically, this effect was restricted to conditions in which participants were free to choose which of two keys to press (compared to conditions with explicit action instructions). It has also been shown that intentional binding can be disrupted by theta-burst transcranial magnetic stimulation (TMS) over the pre-SMA (Moore, Ruge, Wenke, Rothwell, & Haggard, 2010). Curiously, this effect manifested differently between individuals. For some participants, binding of actions to subsequent outcomes was stronger, while for others, binding of outcomes back to actions was stronger.

While the picture that emerges from these collective studies is far from clear, it seems likely that premotor mechanisms underlying volition also contribute to the SoA. This may occur when efference copies of voluntary actions from frontal areas are compared to experienced sensory events, most likely in the parietal cortex.

Further clues to the mechanisms underlying agency may be found in populations with disordered SoA.

1.4.9. Disordered Sense of Agency

In healthy populations, SoA is not always an accurate perception of reality. There is a tendency to attribute frequent or desired outcomes to the self more than objective contingencies dictate. Interestingly, this 'self-serving' bias is diminished in cases of depression (Bandura, 1989; Alloy & Abramson, 1979). One should therefore be careful to distinguish between abnormal and inaccurate SoA when investigating agency in cases of pathology. The use of implicit measures of agency avoids this potential confound.

Delusions of control are a key feature of schizophrenia, in which objectively self-caused actions may be attributed to external sources, possibly arising from a breakdown of comparator mechanisms (Frith, Blakemore, & Wolpert, 2000). Interestingly, schizophrenic patients, despite being more likely to explicitly attribute actions to external sources, show stronger, rather than weaker, intentional binding (Haggard, Martin, Taylor-Clarke, Jeannerod, & Franck, 2003). These results are a useful reminder that *judgements* and *feelings* of agency do not always tell the same story. Interestingly, the observed difference in intentional binding occurs on top of a general tendency to underestimate time perception intervals in schizophrenia, a tendency also seen in PD (Artieda, Pastor, Lacruz, & Obeso, 1992).

The implication of dopaminergic pathways in the experience of volition suggests that PD patients might also show disordered SoA. Indeed it has been shown that manipulation of dopaminergic medication does influence implicit measures of agency (Moore et al., 2010). The intentional binding of PD patients OFF medication (i.e. *hypo*dopaminergic) was not significantly different to healthy controls. Yet patients ON medication (i.e. *hyper*dopaminergic) showed significantly stronger binding than while OFF medication. These results suggest that changes in SoA probably do not change linearly with availability of DA, but that DA manipulation does exert an influence on detection of action-outcome events, and thus SoA. Despite the known involvement of DA in reward processing (Schultz, 2013), there are currently no studies that investigate SoA in PD using experimental contexts that also manipulate reward.

1.4.10. Sense of Agency – conclusions

Most previous laboratory research on SoA has lacked ecological validity. Human actions outside the laboratory are embedded in a rich perceptuomotor, affective and social landscape. Studies of agency have investigated associations between one simple action (Haggard, Clark, & Kalogeras, 2002) or very few actions (Sirigu et al., 1999), and an arbitrary effect, such as a tone, or a very direct effect, such as visual feedback of one's own action. These actions are generally instructed by an experimenter. Outside the laboratory, however, people can generally select among several possible actions in a given situation, and they select a particular action because of a goal they wish to achieve.

Chapters 4, 5, 6 and 7 of this thesis represent a first attempt to use features of reward-based decision-making to study agency in health and disease.

1.5. Overview of experimental work

Chapter 2: Awareness of intentions and actions in Parkinson's disease

Experiments 1 and 2 examine the effect of dopamine manipulation on awareness of actions and intentions in a Libet task. Experiment 1 compares patients with Parkinson's disease both OFF and ON dopaminergic medication, and healthy controls. Experiment 2 compares patients with Parkinson's disease both OFF and ON deep brain stimulation (DBS).

Chapter 3: Neural correlates of intentions and actions in Parkinson's disease

Experiment 3 examines neural activity preceding actions in patients with Parkinson's disease, both off and on DBS, using electroencephalography (EEG). Neural activity is compared to awareness of action and intentions in a Libet task.

Chapter 4: A novel paradigm combining reversal learning and intentional binding

Experiments 4, 5 and 6 describe a novel experimental paradigm to study the Sense of Agency, combining intentional binding with reversal learning.

Chapter 5: Further investigations of agency in reward-based decision-making

Experiments 7 and 8 address critical potential confounds in the novel paradigm introduced in chapter 4.

Chapter 6: A meta-analysis of results from studies employing the novel paradigm

This chapter describes a **multi-study statistical analysis** to determine the reliability and strength of key effects observed in chapters 4 and 5, and from work external to this thesis.

Chapter 7: The Sense of Agency and decision-making in Parkinson's disease

Experiment 9 applies the novel paradigm to a Parkinson's patient group to determine if and how neural features of the pathology are associated with key effects observed in chapters 4, 5 and 6.

Chapter 2: Awareness of actions and intentions in Parkinson's disease.

2.1. Introduction

The following two chapters will describe three experiments using the Libet paradigm (see section 1.3.5.), comparing data from patients with Parkinson's disease (PD) and healthy controls.

The experiments in this chapter were behavioural studies in which the presence of dopamine-related symptoms were manipulated either by medication (experiment 1) or Deep Brain Stimulation (DBS) (experiment 2).

Chapter 2 will describe a second study using DBS (experiment 3) in which electroencephalography (EEG) was used to further elucidate the underlying brain mechanisms of voluntary action in PD.

2.1.1. Actions and intentions in Parkinson's disease.

As outlined in the introduction (section 1.3.9.), PD is a neurological disorder affecting motor control. Motor deficits are associated with irregular dopaminergic (DA) availability from pathways linking the basal ganglia to the Supplementary Motor Area (SMA) (Jahanshahi et el., 1995).

Another feature of PD is irregular time perception. Research seems to indicate that the so-called 'internal clock' may be disturbed in PD. Patients typically underestimate time intervals in verbal estimation and generate overlong intervals in interval reproduction (Pastor, Artieda, Jahanshahi, & Obeso, 1992). Diminished performance in interval judgement tasks has been linked to reduced DA levels (Meck, 2005). There is much less

available data describing timing judgements for punctate events, however there is evidence that awareness of action intentions is also delayed in PD compared to healthy controls (Tabu et al., 2015).

The precise role of DA in disordered time perception in PD is unclear. Studies of patient groups OFF dopaminergic medication have suggested that reduced DA activity does not directly translate into slower movements or delayed awareness of these movements. Rather, it has been suggested (Ganos et al., 2015) that detection thresholds of motor intentions may be shifted due to suboptimal accumulation of evidence for a specific motor action. In PD, evidence accumulation may be disrupted in the presence of motor-related symptoms and alleviated by various methods, including medication and DBS. Experiments in this chapter and following investigated the influence of pathology and DA-related interventions on awareness of actions and intentions in PD.

2.1.2. Aims of this chapter

Experiments 1 and 2 used a Libet task to compare the perceived times of actions and intentions in PD patients while OFF and ON intervention (medication or DBS). In experiment 1, a healthy control group was also used to test for a general effect of pathology.

There were two critical hypotheses. Firstly, if delayed intention awareness is a feature of pathology generally, then patients OFF intervention should show a stronger delay than healthy controls (as in Tabu et al., 2015). Secondly, diminished or irregular DA projections from the basal ganglia to the cortex may be a critical influence on action and intention awareness. If so, patients OFF intervention should show a stronger tendency

to give later timing judgements than the same patients ON intervention. The cohorts and critical comparisons of experiments 1 and 2 are shown below in figure 2-1.

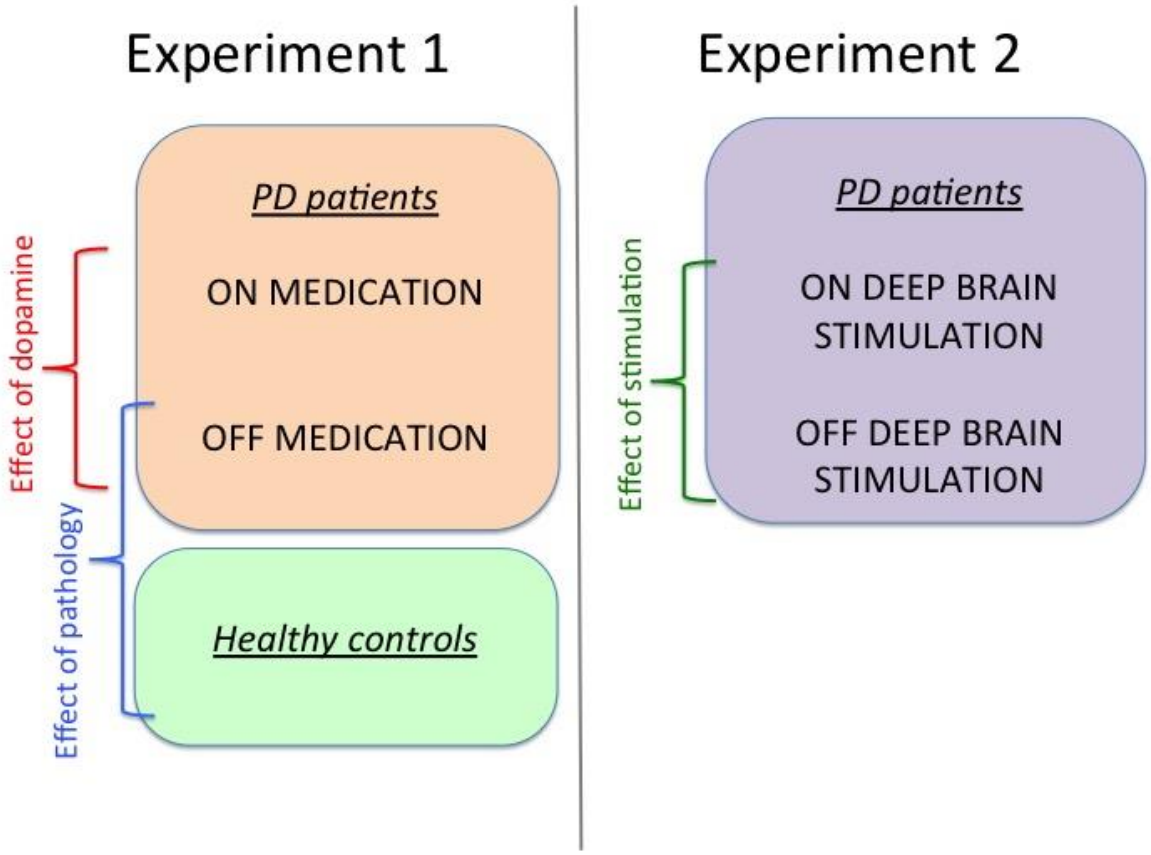

Figure 2-1: Outline of experiments 1 and 2 showing cohorts and critical comparisons.

Previous research on action and intention awareness in PD suggests a number of potential outcomes in experiments 1 and 2. These are represented below in figure 2-2. Assuming a normal readiness potential (RP) prior to movement onset in healthy subjects (panel A of the figure), one may expect that intention awareness but not action awareness will be delayed in patients OFF intervention compared to healthy controls, as in panel B (Tabu et al., 2015). Alternatively, the diminished amplitude in the early premotor RP could result in an early intention awareness followed by late action awareness as in panel C. (Jahanshahi et al., 1995).

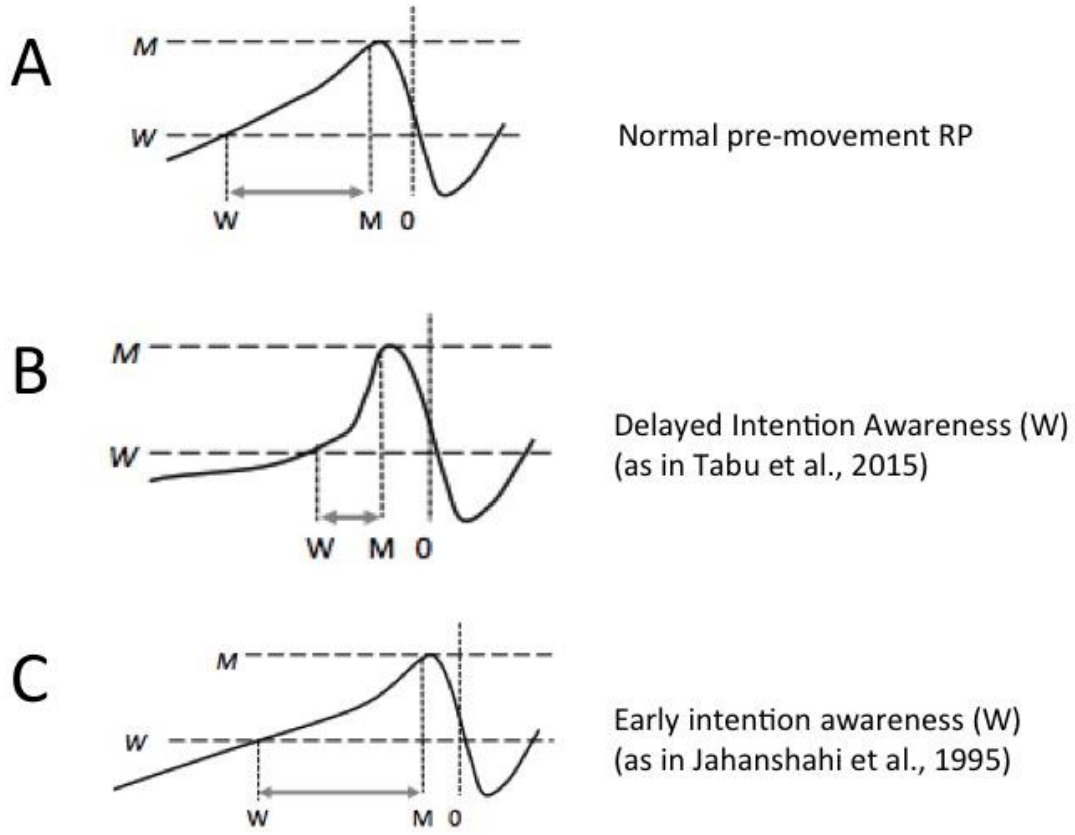

Figure 2-2: Predictions derived from previous research, assuming a normal pre-motor RP in healthy subjects.

2.2. Experiment 1: Intervention manipulation: medication

2.2.1. Introduction

In experiment 1, PD patients and healthy controls completed a simplified version of the Libet task. While fixating a rotating clock, participants pressed a key, and then reported the position of the clock at the time that they made the movement (M judgement) or first felt the urge to move (W judgement).

PD patients completed the task twice. Once while ON their usual prescribed dopaminergic medication and once while OFF medication. Withdrawal of medication as a dopaminergic manipulation has been successfully used in previous research investigating cognitive functions in PD. However, results from these studies show that

patients ON medication may show either benefits or deficits relative to their performance OFF medication (Cools, Barker, Sahakian, & Robbins, 2001; Frank, Seeberger, & O'Reilly, 2004). More information is therefore required to elucidate the effect that dopaminergic medication may have on specific cognitive functions.

2.2.2. Method

Participants

All experiments in this thesis conformed to the Declaration of Helsinki. All studies in chapters 2 and 3 were approved by the Hamburg Ethics Committee. All participants in this experiment agreed to participate in the study and signed a consent form. They all reported normal or corrected to normal vision and hearing.

40 PD patients (27 males, mean age ± SD = 58 ± 8) attending the Department of Neurology of the University Medical Centre Hamburg-Eppendorf agreed to participate and were included in the study. Inclusion criteria were: a diagnosis of PD according to the UK Parkinson's Disease Society Brain Bank criteria, as well as stable treatment and clinical condition for at least 4 weeks prior to the study.

Exclusion criteria were: any major concurrent neurological or psychiatric disorders (no exclusions), a score below 25 on the Montreal Cognitive Assessment (MOCA) (no exclusions) or a score of 16 or more on the Beck Depression Inventory (BDI) (three exclusions). One further participant was excluded due to a self-declared inability to understand the task. We therefore analysed data from 36 patients.

36 healthy, age-matched (11 males, mean age ± SD = 59 ± 10) volunteers in Hamburg were also included in the study and served as a control group. Control participants were

recruited from spouses and friends who accompanied patients to the clinic. One control subject was removed due to a BDI score above 16.

For all patients the following demographic and clinical data were collected: Hoehn & Yahr disease severity score (Hoehn & Yahr, 1967), motor impairment in both OFF and ON medication conditions according to the Unified Parkinson's Disease Rating Scale Part 3 motor examination (UPDRS-III) and Levodopa (L-dopa) and dopamine-agonist (D-Ag) equivalent daily dose (LEDD) (Tomlinson et al., 2010).

The Montreal Cognitive Assessment (MOCA) was employed to assess overall cognitive function including: digit span forward and backward for working memory, phonological word fluency and categorical word fluency for executive functions and language, and Matrix Cancellation Features Target (MCFT) for selective visual attention.

All patients underwent a neuropsychological assessment including the Beck Depression Inventory (BDI) and the Questionnaire for Impulsive-compulsive Behaviour Disorders in Parkinson's disease (QUIP-RS) (Weintraub et al., 2012).

English language versions of all clinical instruments are included for reference in Appendix I.

Procedure

The principle features of the Libet task are described fully in section 1.3.5.

Participants were seated at a standard computer keyboard and screen. They fixated a clock with a single rotating hand. The clock diameter was 20 mm and the hand rotated every 2560 ms. After waiting for the clock to complete one full rotation, and at a time of their choosing, participants pressed the space bar. The clock continued to rotate for a random interval (between 1500 and 2500 ms) and then stopped. According to condition,

participants were then prompted to judge the position of the clock hand either at the moment they pressed the button (M judgement), or at the moment they first experienced the urge to press the button (W judgement). Participants reported their judgements verbally and the experimenter entered the number using the keyboard.

Participants completed 8 counterbalanced blocks of 10 trials. W judgements were elicited in half the blocks and M judgements in half the blocks. At the start of the experiment, participants completed a short training block of each type of trial. This training block was repeated until participants were confident with the task. One participant declined to participate further due to lack of understanding during the training block.

Patients completed the experiment twice. They arrived at the experiment having not taken their prescribed medication for a period of at least 12 hours. The experiment was then run for the first time, after which the patients underwent the clinical assessments detailed above. Patients were then administered medication and waited half an hour before completing the experiment again and undergoing a second clinical assessment.

There is therefore an inevitable order confound, as patients always completed the task OFF medication before ON medication. This is addressed in the discussion (section 2.2.4.)

Control participants completed the experiment and underwent the clinical assessment only once.

All stimuli were presented using LabView 2012 (National Instruments, Austin, TX).

Analysis

The dependent variable was calculated by subtracting the W or M judgement from the actual time of the button press, as captured by the programme. A negative score therefore indicates how long *before* the button press the participants perceived the W or M event to occur.

The dependent variable was then analysed in two critical comparisons:

Firstly, to test for an effect of pathology, the data from healthy participants and patients OFF medication was submitted to a 2x2 between-subjects ANOVA with factors task (W or M judgement) and group (patient or control).

Secondly, to test for an effect of dopamine, data from patients was submitted to a 2x2 repeated-measures ANOVA with factors task (W or M) and state (ON or OFF medication).

2.2.3. Results

Clinical data from patients are displayed below in table 2-1. Clinical data from controls are displayed below in table 2-2.

ID	Hoehn & Yahr	UPDRS OFF	UPDRS ON	LEDD	MOCA	BDI	QUIP
1	2.0	27	17	200	26	14	11
2	2.0	23	9	300	28	8	7
3	1.0	29	13	150	29	1	0
4	2.5	48	33	300	29	6	7
5	2.0	24	17	200	29	5	11
6	1.0	30	21	100	29	9	19
7	2.0	43	28	200	28	10	19
8	2.5	31	21	200	27	7	21
9	2.0	31	19	200	28	7	12
10	2.5	28	19	250	28	3	3
11	3.0	32	24	200	27	3	0
12	2.0	18	14	100	28	6	3
13	2.0	32	23	200	26	6	19
14	2.0	31	20	200	28	4	25
15	2.0	22	12	200	28	11	18
16	2.0	15	10	200	30	13	25
17	2.0	19	13	150	27	5	0
18	2.0	23	15	200	27	8	3
19	2.0	35	23	200	26	5	14
20	2.0	34	23	300	29	3	0
21	2.0	20	14	200	26	1	0
22	1.0	18	13	150	29	12	2
23	2.0	27	20	200	30	4	12
24	2.0	18	13	200	26	10	2
25	2.0	19	12	150	29	2	0
26	2.0	36	22	200	29	3	13
27	2.0	24	12	200	27	7	5
28	2.0	38	25	150	27	3	3
29	3.0	29	20	200	28	8	16
30	1.0	7	4	200	28	12	21
31	2.0	21	15	200	30	1	4
32	1.0	10	4	150	30	6	24
33	3.0	51	34	300	26	4	7
34	1.5	22	16	100	29	13	25
35	2.0	21	15	200	30	6	4
36	1.0	9	5	150	28	6	22
MEAN (±SD)	2 (±0,5)	26 (±10)	17 (±7)	200 (±50)	28 (±1)	6 (±4)	10 (±9)

Table 2-1: PD patients' clinical data from experiment 1.

ID	MOCA	BDI
C1	27	2
C2	29	0
C3	26	0
C4	26	1
C5	29	10
C6	28	4
C7	29	1
C8	26	4
C9	28	1
C10	30	8
C11	30	11
C12	27	5
C13	26	3
C14	30	0
C15	29	0
C16	30	1
C17	28	0
C18	29	2
C19	29	5
C20	28	2
C21	29	8
C22	26	1
C23	29	8
C24	28	2
C25	26	3
C26	30	7
C27	29	11
C28	29	1
C29	27	0
C30	28	14
C31	28	2
C32	29	5
C33	28	0
C34	30	1
C35	30	2
MEAN (±SD)	28 (±1)	4 (±4)

Table 2-2: Healthy controls' clinical data from experiment 1.

A between-subjects ANOVA to compare healthy controls with patients OFF medication yielded a significant main effect of task, with W judgements earlier than M judgements ($F_{(1,69)}$ = 54.99, p < .001, η_p^2 = .44). There was no significant effect of group ($F_{(1,69)}$ < .001, p = *.996*) and no interaction ($F_{(1,69)}$ = .33, p = .57). There was therefore no evidence that PD patients performed any differently to healthy controls.

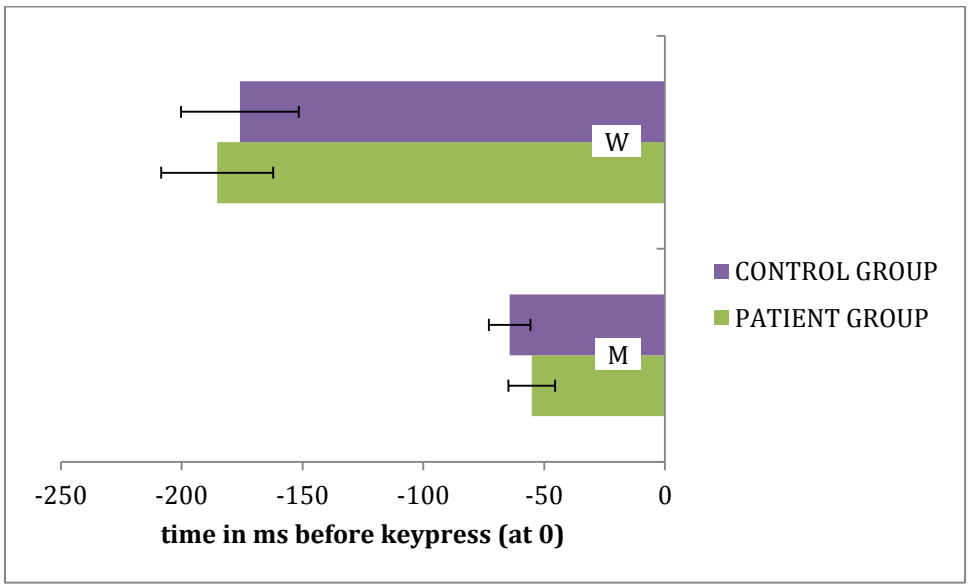

Figure 2-3. W and M judgements of healthy controls and patients OFF medication in experiment 1. Error bars represent standard errors.

Using a within-subjects ANOVA to compare patients OFF and ON medication again yielded a significant main effect of task, with W judgements earlier than M judgements ($F_{(1,35)}$ = 31.2, p < .001, η_p^2 = .47). There was also a significant effect of medication state ($F_{(1,35)}$ = 5.24 , p = .028, η_p^2 = .13) but no interaction ($F_{(1,35)}$ = .62, p = .44), suggesting that patients perceived both W and M events as occurring earlier while on dopaminergic medication.

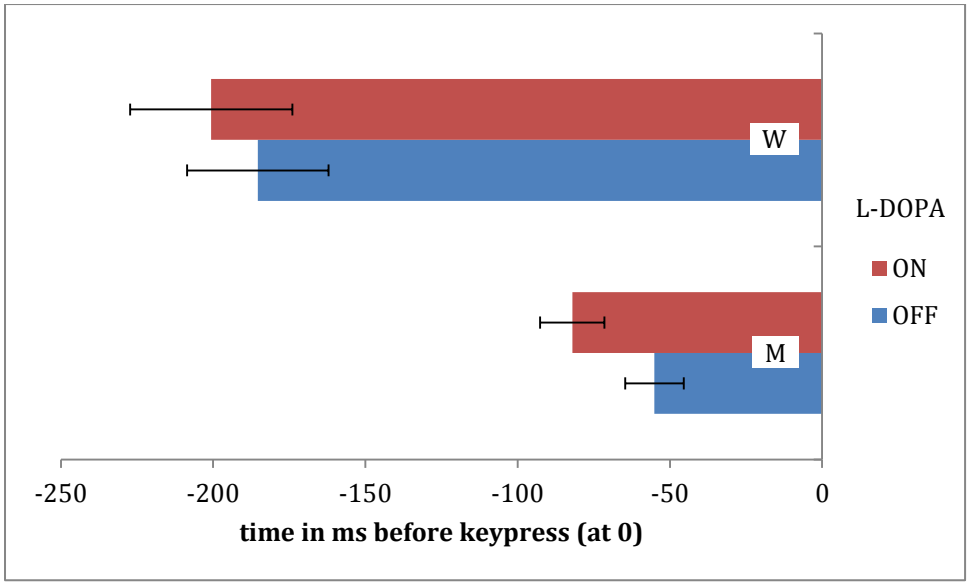

Figure 2-4. W and M judgements of patients while OFF and ON medication in experiment 1. Error bars represent standard errors.

2.2.4. Discussion

The results obtained in experiment 1 show a consistent and strong effect of the task requirement (W or M). Ordinarily, this would only indicate that participants understood and carried out the instructions. However, disordered motor control and time perception are features of PD. One may therefore expect that even if the patient group understood the instructions, their subjective experience of the W and M judgements may differ from healthy controls. The main effect of task, without an interaction, therefore provides evidence that the patient group were able to complete the task with comparable precision to controls.

Furthermore, and contrary to the results obtained by Tabu and colleagues (2015), there was no significant effect of pathology and therefore no evidence that patients OFF medication differed from healthy controls in their subjective perception of action and intentions.

Finally, there was a main effect of dopamine. Patients made significantly earlier W and M judgements while ON medication than while OFF medication. There was no interaction, so it appears that increased DA levels may boost both action *and* intention awareness. These results therefore do not support the predictions illustrated in panels B and C of figure 2-2.

These results contain a necessary order confound as patients always completed the experiment while OFF medication first, and then again while ON medication. Dopaminergic medication takes time to achieve full efficacy, so this confound can only normally be eliminated by running two separate experimental sessions. This was considered inconvenient as participants were selected from a clinical group who mostly lived far from the experimental location. Order-related factors, such as fatigue and boredom cannot therefore be ruled out as alternative explanations for the effect observed.

These results are further considered in the general discussion of this chapter (section 2.4.) with the results from experiment 2.

2.3. Experiment 2: Intervention manipulation: Deep Brain Stimulation

2.3.1. Introduction

Experiment 2 was designed to further test the hypothesis that disordered time perception in PD is linked to irregularity of dopamine projections from the basal ganglia to premotor cortical areas.

Experiment 1 established that PD patients performed with comparable precision to healthy controls in a Libet task. This experiment therefore continues the use of the Libet paradigm without healthy controls. Regularity of the dopaminergic basal-cortical drive was manipulated by means of deep brain stimulation (DBS) of the subthalamic nucleus (STN).

DBS is a surgical therapeutic option for PD patients with intractable tremor and severe cases of dyskinesia (Volkmann, 2004). Surgically implanted electrodes with multiple contacts use high frequency (> 100 Hz) stimulation to provoke a neural disturbance similar to a surgical lesion. Lesions in target areas of the basal ganglia have previously been shown to alleviate symptoms of PD (Koller, Pahwa, Lyons, & Albanese, 1999). The precise mechanisms that result in therapeutic improvement have not been conclusively determined, however it is thought (Santinello et al., 2015) that stimulation leads to a general regularization of neuronal firing patterns in the closed-loop incorporating the basal ganglia, thalamus and cortical motor areas. The effects of DBS on voluntary motor system neurophysiology will be further discussed in chapter 3.

Similar results to experiment 1 were predicted, i.e. patients would report earlier W and M judgements while ON stimulation than while OFF stimulation.

2.3.2. Methods

Participants

All participants agreed to participate in the study and signed a consent form. They all reported normal or corrected to normal vision and hearing, and were paid for their time.

20 PD patients (13 males, mean age ± SD = 65 ± 7) attending the Department of Neurology of the University Medical Center Hamburg-Eppendorf agreed to participate

and were included in the study. Inclusion criteria were: a diagnosis of PD according to the UK Parkinson's Disease Society Brain Bank criteria and a surgically implanted DBS device; treatment and clinical condition stable for at least 4 weeks prior to the study. Exclusion criteria were: any major concurrent neurological or psychiatric disorders, a score < 25 on the Montreal Cognitive Assessment (MOCA) (one exclusion) or a score of 16 or more on the Beck Depression Inventory (BDI) (two exclusions). We therefore analysed data from 17 patients.

Demographic and clinical data were gathered using the same instruments as experiment 1.

Procedure

The general procedure was the same as experiment 1.

Participants completed the experiment twice: once with the DBS device active (ON) and once with the device inactive (OFF). Experimental sessions were followed by clinical assessments as in experiment 1. The required waiting time for the DBS device to change state (active/inactive) was set at half an hour. As the device attains efficacy within this time period, it was possible to counterbalance the order of conditions thereby eliminating the order confound in experiment 1.

Analysis

The dependent variable in the Libet task was again calculated as the judged clock position minus the clock position at the time of the button press.

Data were submitted to a 2x2 repeated measures ANOVA with factors task (W or M) and stimulation state (OFF or ON)

2.3.3. Results

Patients' clinical data are provided below in table 2-3.

ID	Hoehn & Yahr	UPDRS OFF	UPDRS ON	MOCA	BDI	QUIP
D1	3.0	45	30	28	13	18
D2	2.0	49	24	29	1	20
D3	2.0	25	18	28	3	20
D4	2.5	39	24	30	11	6
D5	2.0	27	20	29	6	7
D6	3.0	67	44	28	8	2
D7	2.0	26	16	26	12	10
D8	3.0	46	40	25	4	5
D9	2.0	18	13	30	1	0
D10	2.0	29	14	29	4	19
D11	3.0	18	13	26	12	2
D12	2.0	19	3	26	4	0
D13	2.0	10	1	30	9	4
D14	3.0	42.5	21	28	4	6
D15	2.5	18	13	29	2	0
D16	2.0	38	14	28	2	27
D17	2.0	45	18	29	13	0
MEAN (±SD)	**2 (0.5)**	**33 (15)**	**19 (11)**	**28(2)**	**6(4)**	**9(9)**

Table 2-3: PD patients' clinical data from experiment 2 (DBS)

As in experiment 1, there was a strong effect of task, with W judgements earlier than M judgements ($F_{(1,16)}$ = 18.13, p < .001, η_p^2 = .53). There was no significant effect of state ($F_{(1,16)}$ = 1.66, p = *.22*) and no interaction ($F_{(1,16)}$ = .16, p = .69). There was therefore no evidence that PD patients experienced events differently while ON than while OFF DBS. Note however that the pattern of results is numerically similar to experiment 1 (see figure 2-4 above, and figure 2-5 below).

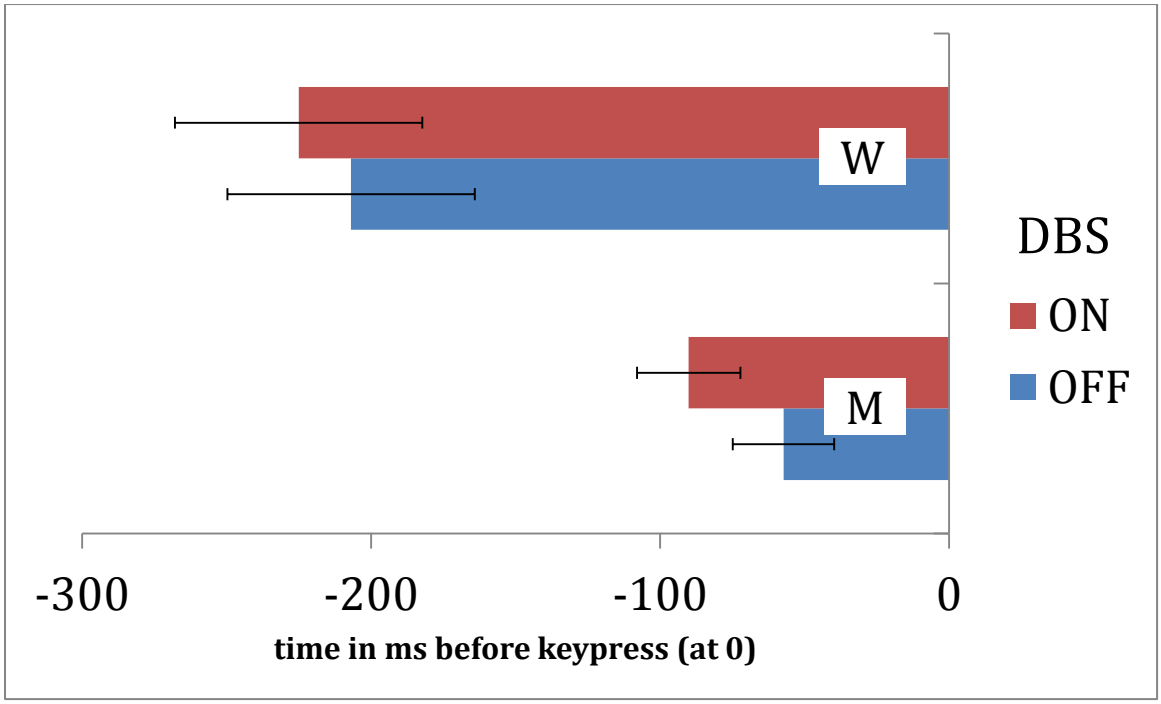

Figure 2-5. W and M judgements of patients OFF and ON DBS in experiment 2. Error bars represent standard errors.

2.3.4. Discussion

Experiment 2 did not find an effect of stimulation state, seemingly at odds with the results obtained in experiment 1. There was a strong effect of task (W/M judgement), indicating that in both states participants were able to complete the experiment with comparable precision.

The lack of an effect of simulation state in experiment 2 may be due to inherent differences in the interventions. DBS alleviates symptoms by influencing the neuronal activity at specific sites in the basal ganglia, specifically the STN, whereas dopaminergic medication has a systemic effect on dopamine receptors. Indeed, it has previously been shown that medication and DBS may have differential effects on PD symptoms (Frank, Samanta, Moustafa, & Sherman, 2007), although the precise mechanisms underlying these differences remain unclear.

Comparison across experiments

Lack of a stimulation effect may also be due to sample size. Experiment 2 comprised roughly half as many participants as experiment 1. Numerically, the pattern of results in both experiments was similar. A follow-up analysis was therefore performed which included results from both experiments. Data were submitted to a 2x2x2 ANOVA with factors task (W or M) and state (OFF or ON), as before, as well as an additional factor of intervention (MED or DBS). A Bonferroni correction for multiple comparisons was applied.

The analysis yielded a significant main effect of task, as expected, with W judgements earlier than M judgements across conditions ($F_{(1, 51)}$ = 45.1, p < .001, η_p^2 = .47) and a significant main effect of state, with participants making earlier judgements across conditions while in an ON state than while in an OFF state ($F_{(1, 51)}$ = 5.99, p = .02, η_p^2 = .11). There was no significant effect of the type of intervention ($F_{(1, 51)}$ = .27, p = .61) and no significant interactions (task x intervention: $F_{(1, 51)}$ = .21, p = .65; state x intervention: $F_{(1, 51)}$ = .05, p = .82; task x state: $F_{(1, 51)}$ = .64, p = .43; task x state x intervention: $F_{(1, 51)}$ = .01, p = .93).

The results of this analysis suggest that there may indeed be a general effect of dopaminergic intervention, independent of delivery method. Experiment 2 might not have detected this effect due to a small sample size. More research is therefore required to delineate the effects and mechanisms particular to each type of intervention.

2.4. General discussion

The results obtained in these experiments suggest that PD patients, despite motor and cognitive impairment, are able to adequately complete a Libet task. Furthermore, there was no evidence that PD patients performed with any diminished precision relative to healthy controls. These results are in contrast to Tabu et al. (2015) who found that awareness of intentions was delayed in PD. This discrepancy in replication may be due to factors that differed between the two studies. As an example, the age of participants may have influenced performance (participants in Tabu et al. were roughly ten years older on average).

It was found that increasing DA availability through medication was associated with earlier judgements of both action and intention awareness in PD patients. Alleviation of DA-related symptoms by DBS did not appear to have the same effect, suggesting that the observed shift in time perception in experiment 1 was due to patients being placed in a *hyper*-dopaminergic state. This may be consistent with the 'dopamine overdose hypothesis' (Vaillancourt, Scholnfeld, Kwak, Bohnen, & Seidler, 2013), according to which administration of levodopa can improve cognitive performance in motor sequence and reversal learning tasks, but increased doses may results in *diminished* performance on these same tasks. More research may be required to determine whether dopamine overload may affect performance in a Libet task, as suggested by the results of experiments 1 and 2.

Interestingly, a follow-up analysis found a significant effect of being ON intervention, which did *not* interact with the specific type of intervention. While medication and DBS implicate different neural mechanisms, they share a common functional impact in

alleviation of PD symptoms. This common impact may indeed be driving earlier judgements in the Libet task, however this was not detected in experiment 2.

The Hallett model of action and intention awareness (see section 1.3.8.) suggests that altered time perception may be due to changes in the timing and rate of motor-related evidence accumulation, or altered thresholds of awareness, or both. The lack of any interaction with the task (W/M) suggests that rate changes are not a convincing explanation, as this would result in steeper or shallower evidence accumulation curves, and therefore manifest as larger or smaller 'gaps' between W and M judgements.

Other possible explanations for the observed effect are illustrated in figure 2-6 (below). The effect of medication suggests that while in a *hyper*-dopaminergic state, RPs may commence earlier, or detection thresholds may be lowered. This could be a direct consequence of increasing the amount of available DA in the neural pathways associated with a voluntary action, resulting in a steeper RP curve (panel B).

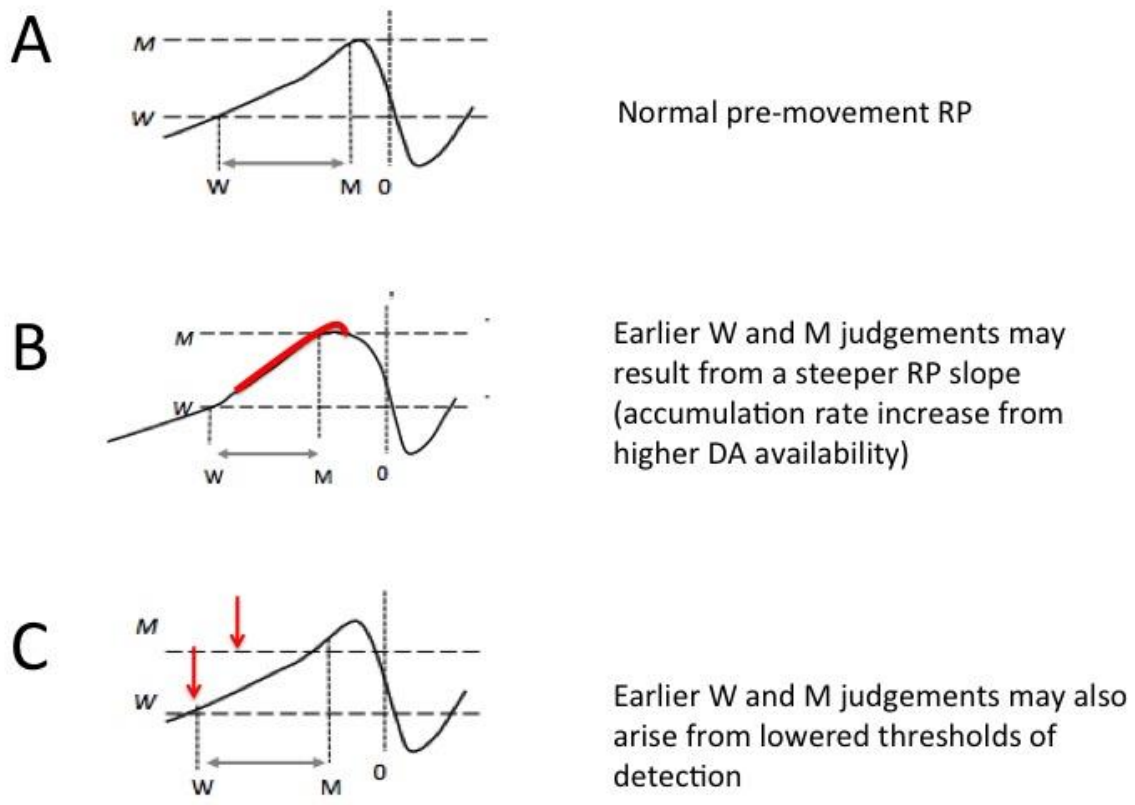

Figure 2-6: Candidate mechanisms for earlier W and M judgements.

Alternatively, one may consider that thresholds of detection are not fixed. If so, the motor-facilitating interventions described here may have lowered the thresholds at which W and M judgements occurred (panel C). Threshold malleability has been described elsewhere. Ganos et al. (2015), for example, provided evidence that thresholds may become more conservative in Tourette's syndrome. In the experiments described here, the alleviation of the motor symptoms of PD may result in a neural environment where the voluntary action signal emerges more readily from the general motor noise. Lowering thresholds would not be problematic, as they would not be prone to inappropriate detection of 'false-alarms'.

The effects of DBS on awareness of actions and intentions on neurophysiology will be further investigated using EEG in the following chapter. The results of experiments 1

and 2 will be considered further in the general discussion of chapter 3, with the results of experiment 3.

Chapter 3: Neural correlates of intentions and actions in Parkinson's disease

3.1. Introduction

The previous chapter examined how PD patients performed in a Libet task compared to healthy controls. It was observed that patients did not display any deficit in completing the task relative to controls. However, patients did perform differently when their dopaminergic symptoms were reduced by means of medication. In a *hyper*-dopaminegic (medicated) state, patients appeared to judge that actions and intentions occurred earlier in time than when in a *hypo*-dopaminergic state (unmedicated). Furthermore, there was some evidence that patients ON intervention judged these events as occurring earlier than when OFF intervention (medication or DBD), regardless of the type of intervention employed.

One theoretical explanation for these changes is that DA availability or the alleviation of DA-related symptoms, influences the accumulation of evidence preceding a motor event, according to the threshold model proposed by Hallett (see section 1.3.8.).

Experiment 3 investigated whether the observed effects could be explained by a shifted onset of the RP, or by a lowering of detection thresholds. This experiment examined neural activity preceding voluntary actions in PD patients, both OFF and ON DBS, using electroencephalography (EEG).

DBS of the sub-thalamic nucleus (STN) in PD patients is associated with faster reaction onset and movement duration. In a positron emission tomography (PET) study, it was reported that motor improvement was accompanied by increased activation in the SMA and premotor cortex (Ceballas-Baumann et al., 1999). Similar results have been

reported for cued free-choice movements (Limousin et al., 1997) where STN stimulation was associated with increased regional cerebral blood flow in the SMA, cingulate cortex and dorsolateral prefrontal cortex. These results demonstrate that regulating communication from the STN to pre-motor areas increases activity in these sites prior to cued reactions or partly voluntary movements. It remains to be seen if similar effects will be observed for actions that are entirely voluntary, as in the Libet task used here.

3.2. Experiment 3: Neural activity preceding a voluntary action in Parkinson's disease: an EEG study.

3.2.1. Introduction

In experiment 3, neural activity measured with EEG was compared to explicit awareness of actions and intentions using the same version of the Libet task used in experiments 1 and 2.

Having established a clear effect of medication, a DBS cohort was used in this experiment to further examine the changes that accompany the alleviation of dopamine-related symptoms by regulation of the dopaminergic basal-cortical drive. Each patient performed the task once with the device active and again with the device inactive. Neural traces were then compared between patients OFF and ON DBS.

Earlier W and M judgements may be the result of an RP that is shifted or 'stretched' away from the moment of action (see figure 2-6). In experiment 2, earlier judgements were observed when DBS was active. One would therefore predict an earlier negative-going component to emerge when DBS is active in the present experiment.

The perceived moments of intention and action may also be perceived as occurring earlier due to a lowering of thresholds of detection. In a less noisy motor environment (i.e when DBS in active), thresholds may be lowered without introducing ambiguity. In this case, earlier W and M judgements could be observed without an early onset RP.

The role of noise in the present experiment will be investigated with an analysis of the variance in RP amplitude in the moments preceding the motor event. There is evidence that prior to a voluntary action, there is a rapid drop off in amplitude variability (Khalighinejad, Schurger, Desantis, Zmirgod, & Haggard, 2017, preprint). This convergence may occur earlier in a less noisy environment, i.e. when DBS is active, thereby reducing ambiguity of motor signals in the late stage of the RP and resulting in earlier W and M judgements. An analysis of convergence analysis will therefore be included in this experiment.

It was predicted that when patients were ON stimulation, earlier W and M judgement would be accompanied by a stronger preparatory drive to frontal motor areas than when patients were OFF stimulation. It was also predicted that increased consistency of cortical activity would be observed prior to action onset in the ON condition relative to the OFF condition.

3.2.2. Method

Participants

All participants agreed to participate to the study and signed a consent form. They all reported normal or corrected to normal vision and hearing.

15 PD patients (13 males, mean age ± SD = 63 ± 8) attending the Department of Neurology of the University Medical Centre Hamburg-Eppendorf agreed to participate and were included in the study. Inclusion criteria were: a diagnosis of PD according to the UK Parkinson's Disease Society Brain Bank criteria, a surgically implanted DBS device and clinical condition stable for at least 4 weeks prior to the study.

Exclusion criteria were: any major concurrent neurological or psychiatric disorders, a score < 25 on the Montreal Cognitive Assessment (MOCA) (no exclusions) or a score of 16 or more on the Beck Depression Inventory (BDI) (no exclusions).

For all patients, demographic and clinical data were collected as in experiments 1 and 2.

Procedure

Participants completed the task in an electrically shielded chamber. Stimuli were presented on a computer monitor with a refresh rate of 60 Hz. The space bar of a standard keyboard was used to input the voluntary action in the Libet task. Participants were fitted with the EEG cap before the experiment began (cap details below) and it was removed at the end of the experiment before debriefing.

The experimental procedure and analysis was the same as experiment 2 (see section 2.3.).

Participants completed the experiment twice: once with the DBS device active and once with the device inactive. The order was counterbalanced across all participants. A period of half an hour was required to elapse whenever the active status of the device was changed, before continuing with the experiment.

EEG activity was recorded referenced to the nose tip from 62 active Ag/AgCl scalp electrodes mounted in an elastic cap with equidistant montage (EASYCAP GmbH, Herrsching, Germany). The electrodes had integrated impedance converters fitted directly into the electrode in order to minimize noise from the surrounding area as well as from movement artefacts.

In order to record EOG activities, two electrodes were placed below the eyes.

Accelerometer signals were collected simultaneously with the EEG activity. The data were bandpass-filtered (0.016–250 Hz) and digitized (sampling rate: 1000 Hz) using BrainAmp amplifiers (BrainProducts, Munich, Germany).

EEG pre-processing

Data from EEG recordings were processed and analysed in Matlab R2015a (MathWorks, MA, USA) using the EEGLab toolbox (Delorme & Makeig, 2004). All signals were downsampled to 250 Hz and low-pass filtered at 30 Hz. Due to the relatively slow fluctuations of the RP, no high pass filter was used. A notch filter at 6 Hz was applied to traces recorded when DBS was active, in order to eliminate a systematic electrical artefact of the DBS device. This artefact can be seen in sample recordings in Appendix II, figure 1.

Data epochs were defined as the period from 2 s prior, to 1 s after the moment of the keypress, as recorded by the programme. Overlap of epochs was impossible, due to the timing of the experimental task.

As mentioned in the introduction (section 1.3.6.), the selection of an appropriate baseline is a critical consideration in EEG experiments measuring RPs. Traditionally, the period from 2500 to 2000 ms before the action event is used, however this involves the

implicit assumption that the RP does not begin until after this time (Shibasaki & Hallett, 2006). This assumption, although standard, is rarely acknowledged and has been questioned by several authors (notably Verbaarschot, Farquhar, & Haselager, 2015; Khalighinejad, Schurger, Desantis, Zmigrod, & Haggard, 2017).

An alternative approach was successfully introduced by Khalighenjad and colleagues (2017) on the premise that the only certain element of an action event is the action itself. They used a baseline period from -5 to +5 ms around the action. This method was validated by running a parallel analysis on demeaned data (using the entire epoch as a baseline), which produced consistent results. This choice of baseline is also critical to the analysis of variance convergence described below, and was consequently used throughout this experiment.

Ocular movement artefacts and other abnormalities were rejected based on visual inspection in an Independent Component Analysis (ICA). Finally, entire data epochs were removed if values exceeded ±150 μv.

Data from 2 participants were excluded following pre-processing due to the limited number of usable trials.

EEG Analysis

Preliminary inspection of the data showed that a characteristic RP (defined by a slow, predominantly negative-going component) was maximal at site FCz. Consequently this electrode site was used in the following analyses. Two variables were calculated for analysis: mean RP amplitude across trials and variability (SD) of amplitude across trials. To compare these variables between OFF and ON states, data from the 2 s preceding the action event were divided into four 500-ms windows. The last time bin was shortened

by 5 ms to avoid overlap with the baseline period. Taking the action event as time point 0, the period of these windows was therefore as follows, in ms: [-2000 to -1500], [-1500 to -1000], [-1000 to -500], and [-500 to -5]. Data were then submitted to four separate 2x2 ANOVAs, one for each time bin, with factors task (W/M) and DBS state (OFF/ON).

Alpha values were corrected for the 4 ANOVAs performed, using a partial Bonferroni correction. Full Bonferroni corrections have been criticised for being too conservative, especially when test outcomes are highly correlated, as here (Perneger, 1998). When correlation information is available, the mean correlation between variables may be used to adjust the full Bonferroni correction (Sankoh, Huque, & Dubey, 1997). Adjusted alpha values were calculated using freely available online interfaces (Uitenbroek, 1997) and are reported below in the relevant results sections.

Across trial variance was further analysed using cluster-based permutation analysis (for guidelines see Maris & Oostenveld, 2007). This method involves a user-defined selection of electrodes and time-window of interest. Each data point therefore represents a channel-time pair. Data points are statistically compared between experimental conditions – in this case DBS OFF/ON – using a t-test. Significant channel-time pairs are then 'clustered' according to temporal and spatial adjacency. Note that here the t-test is used only to determine candidacy in clusters, not as a final statistical value. For each cluster, the t-values of all pairs are summed and compared to a critical value obtained using the Monte Carlo method (data are randomly partitioned a user-defined number of times to create a reference distribution). Importantly, this kind of analysis makes no *a priori* assumptions about electrode sites, time periods or specific ERP components.

For this experiment, the time-window of interest was -2 to 0 s relative to the action. The minimum number of neighbouring electrodes required to form a cluster was 2. The number of random partitions required to create the reference distribution was 1000. Data from W judgement blocks and M judgement blocks were analysed separately.

Note that time-pair data cannot be used to localize activity to specific brain areas. Rather, the data indicate which electrode sources contributed maximally to an overall significant difference in activity across the scalp.

3.2.3. Results

Patients' clinical data are provided below in table 3-1.

ID	Hoehn & Yahr	UPDRS OFF	UPDRS ON	MOCA	BDI
DE1	2.5	45	16	30	11
DE2	2	35	5	missing	15
DE3	2	29	16	29	6
DE4	2	28	10	30	4
DE5	2.5	55.5	44	27	16
DE6	2	31	14.5	27	9
DE7	2	17	6.5	30	5
DE8	3	42.5	21	30	4
DE9	2	37	13	missing	missing
DE10	3	18	7	26	12
DE11	2	missing	18	29	13
DE12	3	34	26.5	25	21
DE13	3	36,5	30	28	12
DE14	2	48	19,5	29	missing
DE15	2.5	55.5	22	28	13
MEAN (±SD)	**2.32 (.4)**	**35 (13.1)**	**17 (10.4)**	**28 (1.7)**	**10.8 (5)**

Table 3-1: PD patients' clinical data from experiment 3

Behavioural results

As in previous experiments there was a strong effect of task, with W judgements perceived as occurring earlier than M judgements ($F_{(1,14)}$ = 10.07, p = .007, η_p^2 = .42). There was no significant effect of state ($F_{(1,14)}$ = .08, p = .78) and no interaction ($F_{(1,14)}$ = .11, p = .75). There was therefore no evidence that PD patients experienced the W and M events differently when ON DBS than when OFF.

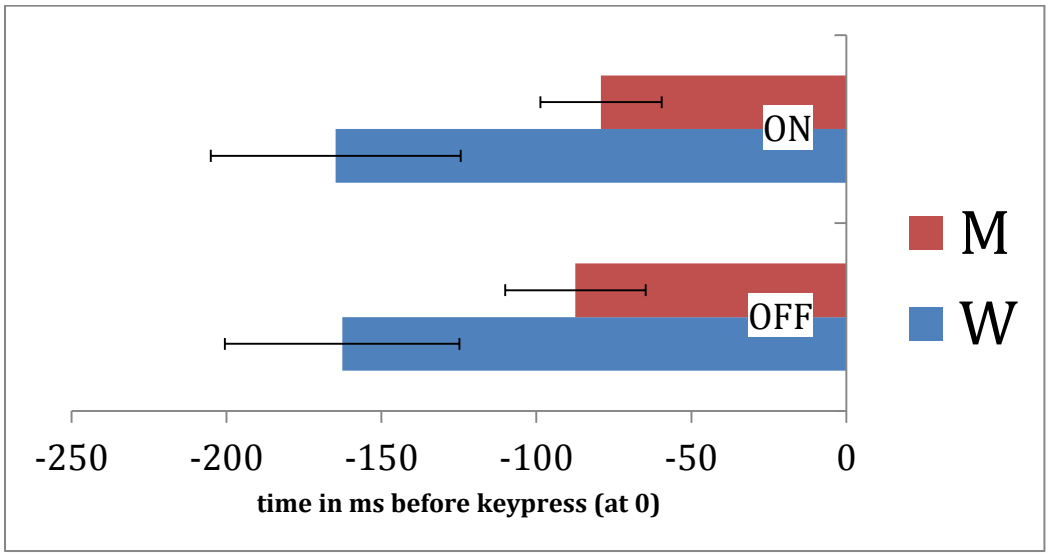

Figure 3-1: W and M judgements in experiment 3.

EEG analysis

Figure 3-2 (below) shows the mean amplitude at FCz before and after the onset of movement separately for W and M judgement data. The use of a baseline from around the time of the action results in a non-traditional view of the RP. Values remain predominantly positive throughout the pre-motor period and are necessarily zero at movement onset. Nevertheless, the data obtained show the characteristic shape of a consistently negative-going shift that defines the RP.

Amplitude was compared across task and state conditions in four time bins within the 2s period immediately prior to the action onset. Full results are displayed below in table 3-2 and figure 3-2. A partial Bonferroni correction for 4 comparisons using correlation adjustment yielded an alpha value of .042. This correction is mild compared to a full Bonferroni correction, reflecting the strong time-dependence of the data. There was no significant main effect of task and no significant interaction in any time bin. This was unsurprising as the motor aspects of the Libet task were identical in both W and M judgement conditions. There was a significant main effect of DBS state in the first time bin ($F_{(1,12)}$ = 6.52, p = .03, η_p^2 = .35) showing that recordings during this time period were significantly stronger when DBS was active than when DBS was inactive.

	BIN 1		**BIN 2**		**BIN 3**		**BIN 4**	
	F(1,12)	*p*	*F(1,12)*	*p*	*F(1,12)*	*p*	*F(1,12)*	*p*
TASK (W/M)	< .01	0.98	0.29	0.6	0.27	0.61	0.1	0.75
STATE (OFF/ON)	**6.52**	**0.03***	3.09	0.1	0.17	0.69	0.8	0.39
INTERACTION	0.09	0.77	0.02	0.89	0.74	0.41	2.08	0.18

Table 3-2: Results of analysis of mean amplitude across trials at site FCz in experiment 3, across four time bins prior to action onset. Each bin represents a period of 500 ms. Alpha value for each comparison was .042 because of adjustment for four correlated comparisons.

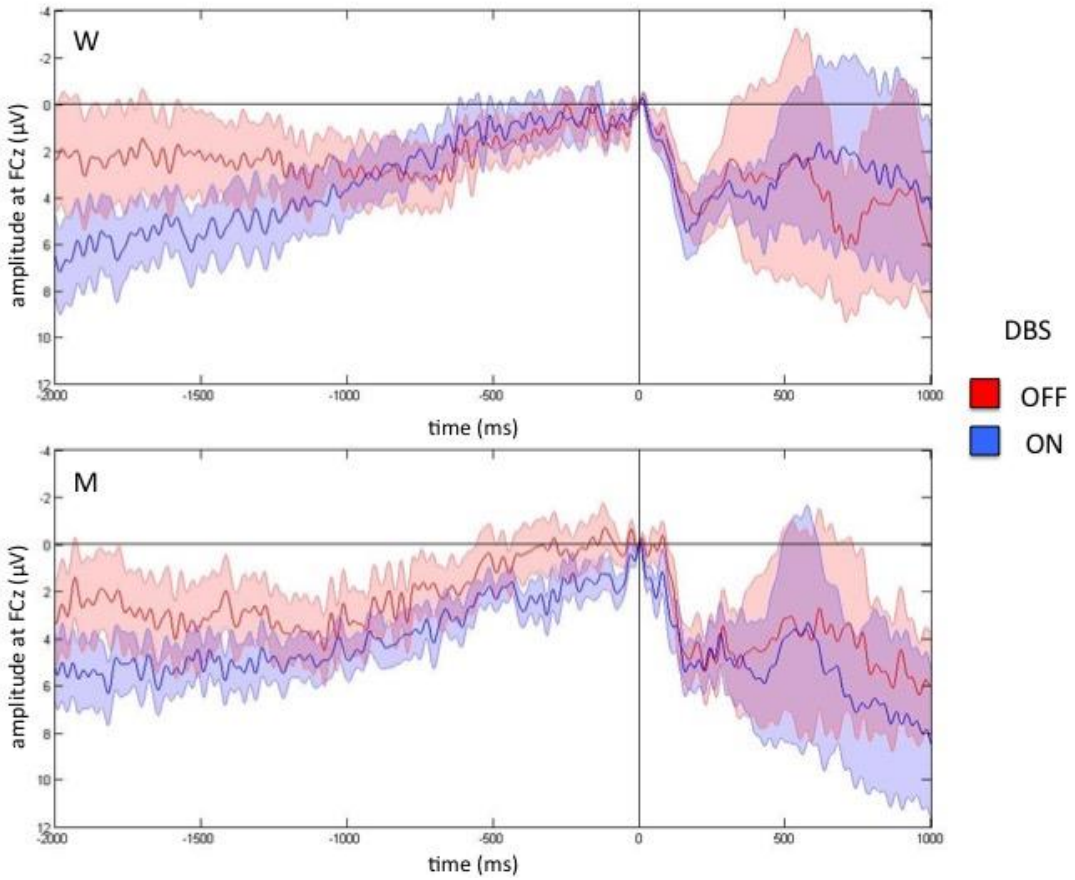

Figure 3-2: Mean amplitude at FCz before and after action (at 0) for W (top panel) and M (lower panel) judgements. Shaded areas represent standard errors.

Across trial SD was similarly compared across task and state conditions in four time bins within the 2s period immediately prior to the action onset. Full results are displayed below in table 3-3 and figure 3-3. A Bonferroni correction for four comparisons using correlation adjustment yielded an alpha value of .02. Again, there was no significant main effect of task and no significant interaction in any time bin. There was a significant main effect of DBS state in the third and fourth time bins (Bin 3: $F_{(1,12)} = 9.38$, $p = .01$, $\eta_p^2 = .44$; Bin 4: $F_{(1,12)} = 11.91$, $p = .005$, $\eta_p^2 = .5$).

	BIN 1		BIN 2		BIN 3		BIN 4	
	F(1,12)	*p*	*F(1,12)*	*p*	*F(1,12)*	*p*	*F(1,12)*	*p*
TASK	1.26	.28	4.02	.07	3.39	.09	3.73	.07
STATE	3.01	.11	4.22	.06	**9.38**	**.01***	**11.91**	**.005***
INTERACTION	1.74	.21	3.1	.1	5.42	.04	3.5	.09

Table 3-3: Results of analysis of SD across trials at site FCz in experiment 3, across four time bins prior to action onset. Each bin represents a period of 500 ms. Alpha value for each comparison was .02 because of adjustment for four correlated comparisons.

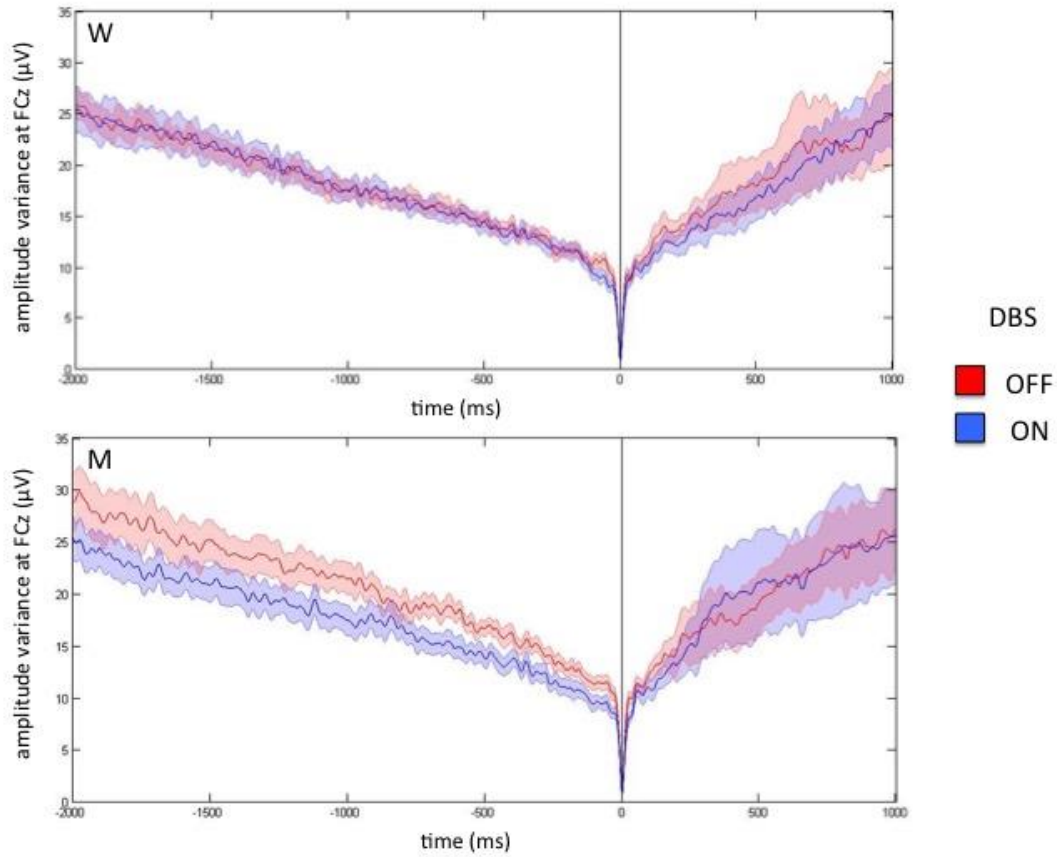

Figure 3-3: Amplitude SD at FCz before and after action (at 0) for W (top panel) and M (lower panel) judgements. Shaded areas represent standard errors.

As the baseline was taken from the 10 ms period around action onset, variance in both conditions at time point 0 was necessarily 0. The observed decrease in variability in figure 3-3 arises partly from this manipulation of the data. However, the results of the statistical comparison indicate that this sudden drop in variability before the voluntary movement was *more marked* when DBS was OFF than DBS was ON.

Cluster-based permutation analysis

No significant clusters were detected for amplitude data in either the W or the M judgement task.

No significant clusters were detected in the across-trial variability data from the W judgement task.

Three significant clusters were detected in the across-trial amplitude variability data from the M judgement task. All three clusters included a broad distribution of electrodes that, along with the inherent features of this type of analysis, prevents any strong assertions about localization of effect (Maris & Oostenveld, 2007). One cluster extended from -690 ms to -380 ms, another from -820 ms to -160 ms and another from -140 ms to -28 ms (see figure 3-4). These time periods are consistent with the results of the ANOVAs described above. Interestingly, the cluster analysis found that this effect was specific to the M judgement task, while the ANOVAs found no significant task specificity.

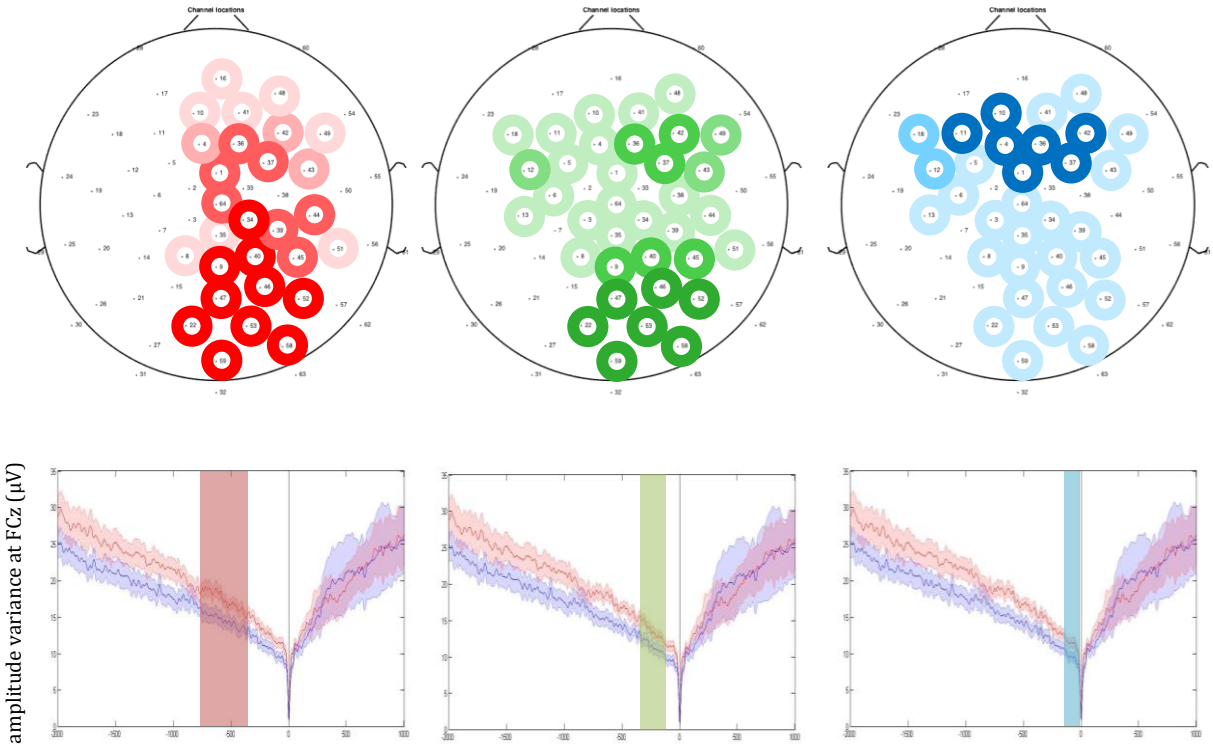

Figure 3-4: Significant clusters detected in permutation analysis. Upper panels: shaded circles mark electrodes included in each cluster. Colour boldness represents the proportion of the total time period of the cluster that individual sites were significant (bolder colours represent a higher proportion of time). For illustrative purposes, the lower panels superimpose the time period of each cluster onto the plotted variability data from FCz. Note however that the permutation analysis included data from all sites.

3.2.4. Discussion

Behavioural data in this experiment did not provide evidence for an effect of DBS on W or M judgements. An effect was predicted based on the results of experiment 1 and the similar functional impact of DBS and dopaminergic medication. These interventions alleviate PD symptoms via different mechanisms and this difference may account for the lack of a detectable effect in the DBS cohort. However, the neurological data presented below suggest that DBS does influence pre-movement neural activity, so alterations in time perception while ON stimulation still seem plausible. Further experimental evidence from larger cohorts is recommended to adequately test this hypothesis.

Comparison of the EEG recordings at site FCz showed that when DBS was ON, RP amplitude was significantly stronger, but only in the first premotor time bin. This early

difference in RP amplitude implies that the use of a baseline from around the time of action was indeed justified in this case. This suggests a disparity of RP onset, followed by a later convergence. Specifically, although RP amplitude was more negative in the OFF state in the early stage, the negative-going shift does not emerge clearly until around 1 s before the action. RP amplitude in the ON state already showed a characteristic negative going trend at 2 s before the action.

The more negative amplitude generally in the early stage of the RP in the OFF condition may represent the effect of increased neural noise. This increased noise may have delayed or masked the characteristic negative-going phase until an unambiguous motor intention emerged. In the ON condition, this motor intention would have been detected earlier. These results therefore provide evidence for a candidate mechanism (RP onset shift) underlying earlier W and M judgements when DBS is active, as seen in experiment 2. However this perceptual effect was not replicated in the behavioural data of *this* experiment, and the convergence of RPs occurred *before* the average time of W and M judgements. Further research is therefore required to determine the robustness of the effect and the precise timing of the implicated neural events.

The analysis of across-trial variability showed that for both the W and M judgement tasks, the usual convergence of amplitude variability in the moments preceding a voluntary action was observed both in the OFF and ON conditions. Interestingly, in the M judgement task the late drop off in variance was significantly more marked in the OFF condition, due to variance remaining significantly higher during the very late stage of pre-motor activity. This delayed convergence may reflect a higher amount of noise generally when DBS is inactive, which is likely associated with the movement-related symptoms of PD.

It is tempting to suggest that the degree of convergence may be used as an index of a consistent precursor process of preparation for voluntary action. The earlier convergence seen in the ON condition might therefore seem to be a candidate explanation for the earlier W and M judgements observed in experiments 1 and 2. However there are two problems with this approach. Firstly, no effect of DBS was found in this experiment, so any association between EEG recordings here and the behavioural results in experiments 1 and 2 remains tenuous. Secondly, the difference in convergence was only observed in the M task in this experiment, so it cannot account for earlier W judgements.

It remains unclear why the convergence difference was only found in the M judgement. One possible explanation is that the W judgement requires a greater degree of cognitive engagement: unlike for the M judgement, there is no physical change that can be sensed by the patient when considering the onset of intention. It has been shown that increased cognitive effort is associated with a reduction in variability (Manohar et al., 2015). In the present experiment, this effect may have driven reduced variability across conditions in the W task, masking or overriding any effect of stimulation. Indeed, the values displayed in figure 3-3 show that variability was remarkably consistent except in the M task when DBS was active. In this particular condition, cognitive effort was low, and variability was reduced only when DBS was active, possibly by alleviation of motor symptoms through a regulated basal-cortical drive.

Taken together, the results of this experiment describe clear differences in neural activity in the moments preceding a voluntary action, depending on the presence or absence of deep brain stimulation. The characteristic negative phase of the RP appeared earlier and showed greater change in scalp potential when DBS was ON.

In the two motor bins immediately preceding the action, a characteristic reduction in variance occurred earlier when DBS was active than when inactive. In a cluster-based permutation analysis, this difference in variance was shown to be significantly different between conditions, across a broad range of central sites immediately before the onset of the action. These shifts in the shape of the RP were considered a candidate mechanism for the earlier W and M judgements observed when DBS was active in the previous chapter. However, this behavioural effect was not replicated in this sample, so this suggestion was rejected. Nevertheless, it is clear that using DBS to regulate the basal-cortical drive influenced the shape and behaviour of the characteristic RP during a Libet task. The behavioural manifestations of these changes remain to be more adequately determined in future research.

The combined results of chapters 2 and 3 are further discussed below.

3.3. Discussion of Chapters 2 & 3

Chapter 2 showed that PD patients were able to complete a Libet task with performance and precision comparable to that of healthy controls. Interestingly, when patients were ON dopaminergic medication, they gave significantly earlier judgements of the perceived time of the intention to act, and of the action itself, than when OFF medication. A cross-experiment analysis also provided evidence that W and M judgements were both significantly earlier when patients were ON an intervention, regardless of what that intervention was.

Chapter 3 attempted to detect the neural activity underlying this effect through an analysis of EEG recordings taken while patients completed the Libet task both OFF and ON DBS. The behavioural effect seen in chapter 2 did not replicate. However, activity

recorded at electrode site FCz and broadly across the central area showed that when patients were ON DBS, the characteristic negative phase of the RP was lengthened and showed a greater shift compared to RPs recorded when patients were OFF DBS. Furthermore, convergence of amplitude variability across trials showed an earlier drop-off in the moments immediately preceding action onset.

An earlier onset to the RP and reduced noise in the environment both predict earlier awareness of action and intention in a threshold model of action preparation (Hallett, 2007). In this sense, it is tempting to conclude that the observed changes to the RP observed here drive the shift in perception described in experiments 1 and 2. However, RP change and earlier action awareness were not observed in the same cohort, so interpretation must be cautious.

There was no healthy control group in experiment 2 or experiment 3 (see figure 2-1). Consequently there is no way to know whether it was the *presence* or *absence* of DBS that drove changes in the RP. It can only be stated that there was a difference between the two states. Furthermore, the limited availability of DBS patients and the time required to complete the task – particularly when EEG was included – were factors that likely resulted in a lack of statistical power in experiments 2 and 3.

The results obtained in these two chapters are consistent with a trend in the literature showing that disturbances of voluntary movement are accompanied by a delay in the awareness of both intentions and actions (Tabu et al., 2015; Moretto et al., 2011). The experiments described here provide further evidence from patients in both hypo- and hyper-dopaminergic states, or with motor symptoms alleviated or not by DBS. There was clear evidence from medication manipulations that the perceptual shift may be driven by dopamine availability, however the results were less conclusive for DBS.

The use of EEG in experiment 3 provided a valuable insight into the patterns of neural behaviour that precede a voluntary action in the presence or absence of DBS. Specifically, it was demonstrated that activation of DBS was associated with shape changes in the characteristic RP. These changes did not cause changes in action awareness. However, the activity observed does match predictions made by threshold models of motor preparation.

Taken together, these results demonstrate that PD and its associated treatments continue to provide a unique and useful opportunity to investigate the neural mechanisms underlying voluntary action. Voluntary actions in PD will be further investigated in experiment 9, which focuses on the sense of agency (see section 7.2.). However it will first be necessary to introduce a novel paradigm that will be used to investigate the sense of agency in the following chapters.

Chapter 4: A novel paradigm combining reversal learning and intentional binding

4.1 Introduction

4.1.1. Limitations of previous Intentional Binding research

Previous laboratory research on Sense of Agency (SoA) using Intentional Binding has often lacked ecological validity. Studies have typically investigated associations between a single action and a single outcome, without any significance or value for the participant (Haggard, Clark, & Kalogeras, 2002).

Outside the laboratory, however, actions are embedded in a rich perceptual, affective and social landscape. People frequently select one action from several possible in a given situation, to achieve a desired goal. However, only a few studies have attempted to link sense of agency with outcome valence. In Takahata et al. (2012), participants' actions caused tones that were associated with monetary rewards or penalties. They found reduced binding for penalty trials compared to neutral or rewarded trials. Similarly, Yoshie and Haggard (2013) found that negative sound outcomes were associated with a reduction in binding compared to neutral and positive sound outcomes. Neither study explored the effects of contingency between participants' actions and the rewards received. When action choices lead to valenced outcomes, there is the possibility for learning, to maximise the benefit arising from one's actions. The nature and strength of action-outcome associations formed in a learning context may differ strongly than associations formed in a more repetitive task. To investigate the intentional binding effect in a dynamic choice environment, this chapter introduces a novel paradigm that draws from research in goal-based decision-making.

4.1.2. Reversal Learning: an established paradigm for investigating goal-based decision-making.

The novel paradigm presented in this thesis makes use of a probabilistic reversal learning approach (Rolls, 2000; Cools, Clark, Owen, & Robbins, 2002), which requires participants to continuously learn and update their action choices to align with changing action-outcome mappings. Typically, participants select between two options, one of which is more frequently rewarded. Periodically, the more favourable option and the less favourable option are inverted without the participant being made explicitly aware. To maximize rewards, participants must therefore monitor the outcome linked to each action, and then correctly update their expectations so as to select their *next* action accordingly (Sutton & Barto, 1991).

A central issue in research on learning is how behaviour changes trial by trial in response to feedback (Daw, 2011). The outcomes in the paradigm described above are valenced: rewards are a positive outcome while non-rewards (or penalties) are negative. It was predicted that the valence of action outcomes might influence not only the intentional binding associated with a given outcome, but also the intentional binding reported on the subsequent trial.

4.1.3. The need for a new paradigm.

Despite the seemingly intuitive connection between the sense of agency and goal-based decision-making, these two research streams have largely evolved in separate experimental disciplines. Action *intention* and action *selection* have rarely been investigated together and there have seemingly been no experimental paradigms designed to study these associated experiences within a single task.

4.2. Experiment 4: Intentional binding and reward-based decision-making.

4.2.1. Introduction

This experiment represents a preliminary and exploratory attempt to study the sense of agency within a goal-based decision-making context.

The 'Libet clock' method was used to measure Intentional Binding (see section 1.3.5. for full description). This method has proven itself robust to replication (Moore & Obhi, 2012) and allows for precise timing judgements to be made without any extensive training on the part of the participant.

The voluntary action event and the sensory outcome of the Intentional Binding task were used as the essential elements in the probabilistic learning paradigm (Rolls, 2000; Cools, Clark, Owen, & Robbins, 2002). The voluntary action was made into a choice point, where one of two actions could be selected. The sensory outcome, either of two different tones (high or low frequency), indicated whether or not that action was rewarded. Further details are provided below.

4.2.2. Methods

All experiments in chapter 4 and 5 were approved by the UCL Research Ethics Committee.

Participants

Sixteen participants (9 female, all right-handed, mean age = 23 years, age range = 18-41 years) completed the experiment and were paid £8/hour plus a bonus for correct

responses. Data from one participant was lost due to a technical error that corrupted a data file. All participants reported normal or corrected to normal vision and hearing.

Procedure

As in chapters 2 and 3, participants were seated at a standard computer keyboard and screen. They fixated a clock with a single rotating hand. The clock diameter was 20 mm and the hand rotated every 2560 ms.

In baseline conditions, participants pressed a key at a time of their free choice, or heard an auditory tone at a random time. In 'operant' conditions, participants both pressed a key *and* heard a tone. The tone occurred 250 ms after the key press. Participants were instructed to wait for one full rotation of the clock before pressing the key. Tones were either high (2000 Hz) or low (500 Hz) in frequency, and lasted 100 ms. The high tone was always the 'correct' tone, and was associated with a monetary reward. The 'F' and 'J' keys of a standard keyboard were used for left and right hand responses.

Following the tone (or the key press if no tone), the clock hand continued to rotate for a random interval between 1100 ms and 2800 ms, and then disappeared. In one block, participants reported the time that they pressed the button. In another block, they reported the time that they heard the tone. They used the keyboard to enter their judgement.

Baseline action and outcome measures were first taken in six separate blocks in pseudorandom order. In three action baselines blocks, participants pressed the key with their left, right or freely chosen hand, according to block. No tone was played. Participants then estimated the time they pressed the key. In three outcome baseline

blocks, participants did not press a key, but heard a high, low or randomly mixed (high or low) tone, according to block. They then estimated the time that the tone was played.

Next, participants completed two counterbalanced 'operant' blocks. In each operant block, one key delivered rewarded high tones with a probability of 0.8 and the other key with a probability of 0.2. The mapping was maintained across a run of trials, until the participant had selected the more rewarding key between five and seven times consecutively (randomized). Probability mappings then reversed. Nine such reversals occurred in each block, so each block involved ten 'runs' of responses. The actual number of trials per block depended on how rapidly each participant learned the 'correct' key.

Participants were explicitly told that the two keys had different probabilities of reward but were not told the exact probability values. They were informed that neither key would ever be a guaranteed win or loss (i.e. probability of reward was always less than 1 and greater than 0). Participants were also told that the probabilities would occasionally reverse, but were not told how often or under what circumstances.

The cumulative total of rewarded trials was displayed at the end of each trial. At the end of each block, all participants were told they had reached the threshold number of rewarded trials required to trigger a bonus. In fact, this threshold was fictitious, and a bonus of £3 for each block was paid at the end of the experiment. This arrangement ensured that participants were not overpaid for prolonging the experiment by repeatedly making incorrect responses.

At the end of each trial, a visual feedback indicating either reward (tick) or non-reward (cross) was presented after each judgement. The visual signal recapitulated the

information previously conveyed by the auditory tone, but was included to facilitate decision-making on the next trial, without placing strong demands on memory.

All six baseline conditions were then repeated in reverse order.

Analysis of baseline trials

No significant differences were observed in the perceived times of key presses in milliseconds for left and right hand responses, forced- or free-choice, or for pre- or post-experiment blocks measures ($p > .05$ for all comparisons). Consequently, all action baseline blocks were collapsed in further analysis.

No significant differences were observed in the perceived times of high and low frequency auditory tones, for mixed or repeated presentation or for pre- or post-test measures ($p > .05$ for all comparisons). Consequently, these were also collapsed in further analysis.

Analysis of operant trials

Perceptual shifts were calculated for each participant and each condition by subtracting the relevant mean baseline error for actions or tones from that in operant trials. A positive action binding measure therefore corresponds to a shift of the perceived time of the action towards its outcome and a negative outcome binding measure to a shift of the perceived time of the outcome towards the action.

Operant trials were categorized according to two design factors: 1. whether the outcome received on the current trial was rewarded (high tone) or not rewarded (low tone), 2. whether feedback on the *previous* trial was rewarded or not rewarded.

4.2.3. Results

Performance

The overall ratio of trials with non-rewarded outcomes to rewarded outcomes was 0.6:1 (mean number of trials per block = 109, SD = 35).

Participants learned the action-outcome contingencies. The criterion for advancement was set at five to seven presses of the more rewarded key, so performance was necessarily 100% before reversal of action-outcome mappings. Reversal events unsurprisingly triggered errors. The proportion of correct choices was submitted to a repeated-measure ANOVA with trial number after reversal as a factor (up to 5 trials). The trial number had a significant effect on participants' performance ($F_{(4,56)}$ = 66.2, p < .001, η_p^2 = .250), see figure 4-1. As the figure shows, participants adapted their responses after a few reversal-induced errors occurred.

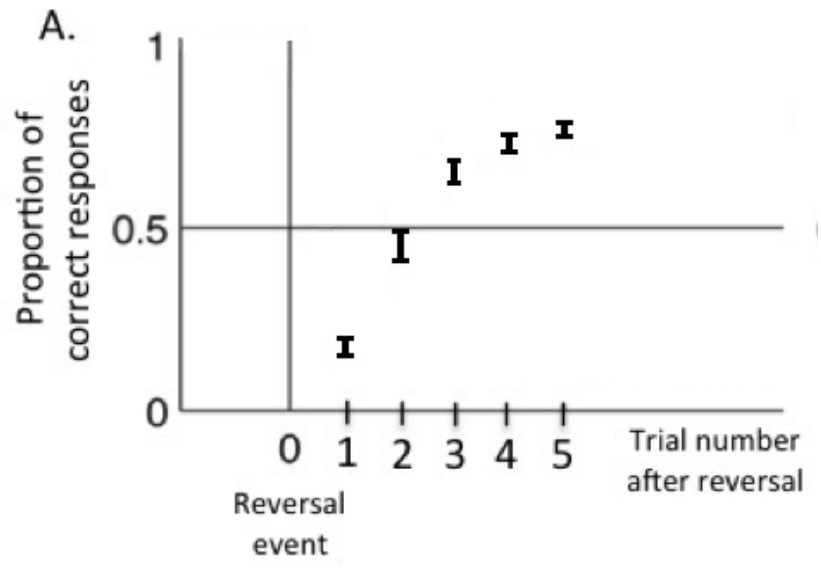

Figure 4-1. Proportion of correct responses after a reversal event in experiment 4. Reproduced from Di Costa, Théro, Chambon, & Haggard (2017)

Intentional binding

Action binding data were submitted to a 2x2 ANOVA with factors of current outcome valence (rewarded or not rewarded), and previous outcome valence. There was a highly significant effect of previous outcome valence ($F_{(1,14)}$ = 9.2, p = .009, η_p^2 = .397), with stronger action binding following a non-rewarded outcome than following a rewarded outcome. There was no effect of current outcome valence ($F_{(1,14)}$ = 1.72, p = .21, η_p^2 = .110), and no interaction ($F_{(1,14)}$ = .01, p > 0.25, η_p^2 = .000).

A similar ANOVA was performed for outcome binding. This showed a significant effect of current outcome valence ($F_{(1,14)}$ = 6.32, p = .025, η_p^2 = .311), with non-rewarded outcomes being more strongly bound towards actions than rewarded outcomes. There was no effect of previous outcome valence ($F_{(1,14)}$ = .02, p = .89, η_p^2 = .002), and no interaction ($F_{(1,14)}$ = 1.89, p = .19, η_p^2 = .119).

Action and outcome binding data are shown in figure 4-2.

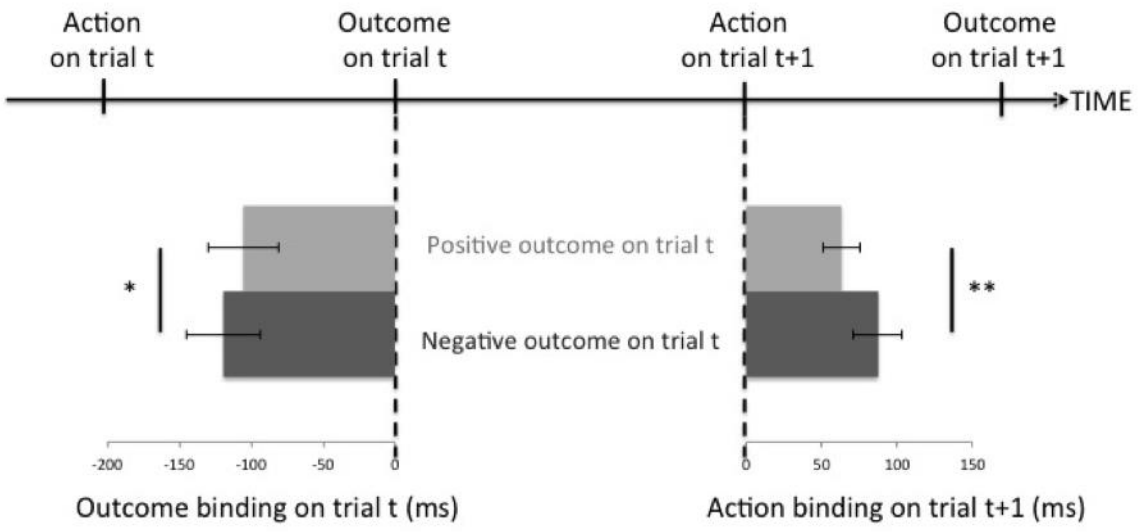

Figure 4-2. Outcome and action binding in experiment 4. Error bars represent standard errors. Reproduced from Di Costa, Théro, Chambon, & Haggard (2017)

4.2.4. Discussion

In a reversal-learning task, non-rewarded outcomes were more strongly bound back to their actions than rewarded outcomes. These results therefore differ markedly from previous studies of binding and valence (Takahata et al., 2012; Yoshie & Haggard, 2013), in which non-rewarded outcomes showed less binding than rewarded or neutral outcomes. This difference may reflect the presence of learning and selection in reward-based decision-making in the current paradigm, which was absent in those previous studies.

Further, action binding on the trial *following* a non-rewarded outcome was stronger than following a rewarded outcome. This is the first time that previous trial outcome is reported to have a *sequential* effect on action binding.

Participants often – but not always – switched keys after an error. It was therefore important to establish that the observed post-error effect from an effect of performing a novel action. A simple linear regression model was fitted to the action binding data, with the following factors: whether or not participants switched keys (β1), whether or not the previous trial was rewarded or not (β2) and the interaction (β4), as well as an error term.

Action Binding = β3 * Interaction + β1 * Switch + β2 * PreviousReward + β4

The switch and interaction factors (β1 and β3) were both non-significant, while previous trial reward was significant (p = .02). These results indicate that the observed

difference in action binding is more likely explained by the previous trial feedback, than by whether or not the participant switched keys.

Engbert and Wohlschläger (2007) showed that action binding reflects learning and prediction processes. The following experiment therefore aimed to replicate this post-error boost of action binding and demonstrate that it is strongest in contexts where learning can be used to maximize reward.

4.3. Experiment 5: Influence of more or less ambiguous learning environment

4.3.1. Introduction

The previous experiment revealed two effects of interest. Firstly, non-rewarded outcomes were more strongly bound to their actions. This effect was in contrast to predictions made from previous work. This experiment therefore aims to replicate the effect.

Secondly, actions on trials following non-rewarded outcomes were more strongly bound to their outcomes. This experiment aims to replicate this effect and investigate if this sequential influence is associated with participants' ability to learn from outcomes.

In this experiment, the novel paradigm described in experiment 4 was used again. However the reward contingency of the two options was further manipulated. In different conditions, it was made more or less easy for participants to use outcomes to inform their action selection. This was achieved by altering the reliability of an outcome given a particular choice. If the effects observed in the previous experiment were indeed associated with learning, binding should be stronger in the easier condition of this

experiment, in which actions predict outcomes more reliably than in the more difficult condition.

4.3.2. Methods

Participants

Sixteen participants (11 female, all right-handed, mean age = 22 years, age range = 18-48 years) completed the experiment and were paid £7.50/hour plus a bonus for correct responses. All participants reported normal or corrected to normal vision and hearing.

Procedure

The design was based on the paradigm described in experiment 4, with some adjustments. Participants completed two blocks of trials (easy/hard) for each type of judgement (action/outcome) for a total of four blocks. In 'easy' blocks, one key delivered rewards with a probability of 0.9 and the other with a probability of 0.1. In 'hard' blocks, one key delivered rewards with a probability of 0.7 and the other with a probability of 0.3.

Analysis of baseline trials

No significant differences were observed in the perceived times of key presses in milliseconds for left and right hand responses, forced- or free-choice, or for pre- or post-experiment blocks measures ($p > .05$ for all comparisons). Consequently, all action baseline blocks were collapsed in further analysis.

Analysis of operant trials

Perceptual shifts were calculated as in experiment 4. In addition to the factors of previous and current trial outcome, the factor of reward contingency was included (easy/hard).

4.3.3. Results

Performance

As in experiment 4, participants learned the action-outcome contingencies (see figure 4-3). Reversal events triggered errors and participants then adapted their responses after a few reversal-induced errors occurred and quickly returned to a plateau of sustained performance. It is worth noting that this plateau was at approximately 90% for the easy condition and 70% for the difficult condition, in line with the probability matching law (Vulkan, 2000).

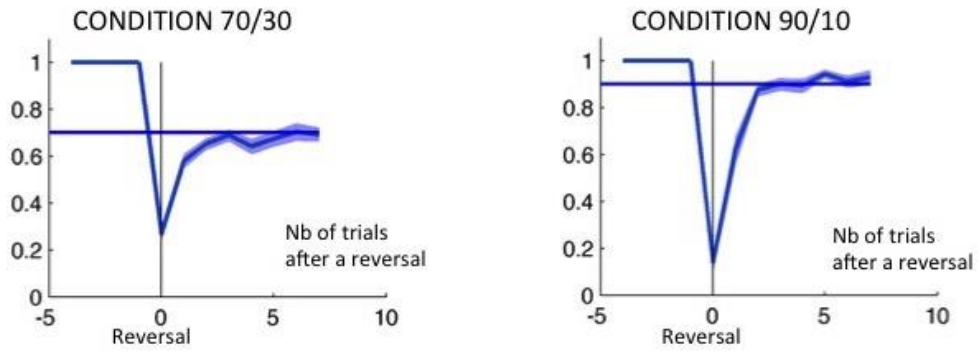

Figure 4-3. Proportion of correct responses before and after a reversal event in experiment 5.

Action binding data were submitted to a 2x2 ANOVA with factors of current outcome valence (rewarded or non-rewarded), previous outcome valence and reward contingency (easy/hard). There was no significant effect of previous outcome valence

($F_{(1,15)}$ = 1.77, p = .203, η_p^2 = .105), no effect of current outcome valence ($F_{(1,15)}$ = .188, p = .67, η_p^2 = .11), and no effect of reward contingency ($F_{(1,15)}$ = 3.54, p = .079, η_p^2 = .191). All interactions were non-significant (p > .05).

A similar analysis was performed for tone binding. Data were submitted to a 2x2 ANOVA with factors as above. There was no significant effect of previous outcome valence ($F_{(1,15)}$ = .019, p = .893, η_p^2 = .001), no effect of current outcome valence ($F_{(1,15)}$ = 2.7, p = .121, η_p^2 = .153), and no effect of reward contingency ($F_{(1,15)}$ = .061, p = .809, η_p^2 = .004). All interactions were non-significant (p > .05).

To determine if the effects observed in the previous experiment were replicated across the experiment, binding results were directly compared across difficulty conditions. Action binding was tested through paired t-tests on the pooled conditions. The previous outcome valence had a significant effect on action binding ($t_{(15)}$ = 2.28, p = .038), with stronger action binding following a non-rewarded outcome than following a rewarded outcome.

Data were also pooled across reward contingency for outcome binding. The current outcome valence had a significant effect on outcome binding ($t_{(15)}$ = -2.40, p = .03), with non-rewarded outcomes being more strongly bound towards actions than rewarded outcomes.

Action and outcome binding data from pooled results are shown in figure 4-4.

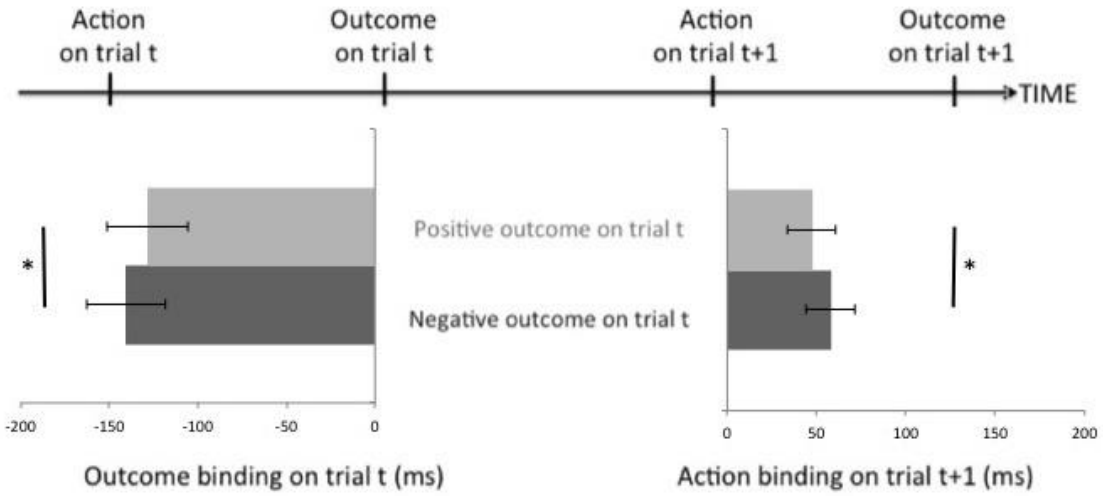

Figure 4-4. Outcome and action binding in experiment 5. Error bars represent standard errors.

4.3.4. Discussion

The purpose of experiment 5 was to examine the effect of a more or less reliable reward contingency on SoA. The experiment also attempted to replicate two effects from experiment 4: an effect of previous trial outcome valence on action binding, and an effect of current trial outcome valence on outcome binding.

When data were separated by reward contingency the effects of outcome valence on outcome binding (current trial) and action binding (subsequent trial) did not replicate. However, an analysis of the pooled data did reveal significant effects, consistent with those observed in experiment 4.

There was no significant effect of reward contingency on action or tone binding. Implicit measures of binding were therefore insensitive to the perceived reliability of action alternatives in this experiment. One explanation for the lack of effect may be that the

levels of the difficulty manipulation were not distinct enough. A third experiment therefore compared intentional binding in *learning* and *non-learning* conditions.

4.4. Experiment 6: Influence of a learning context

4.4.1. Introduction

This experiment was again based on the paradigm introduced in experiment 4.

In experiment 6, the degree of difficulty was not manipulated, but a 'non-learning' condition was included. In this condition, participants made actions and received outcomes as before, but action-outcome mappings were now entirely unpredictable. Participants were explicitly told about the nature of these two conditions. Stronger action binding was predicted in the learning condition compared to the random condition.

4.4.2. Method

Participants

Thirty participants (21 females, all right-handed, mean age = 28 years, age range = 21-53 years) completed the experiment and were paid £7.5/hour plus a bonus for correct responses and precision. The number of participants was increased, compared to experiment 4, to allow us to correlate intentional binding measures with learning measures across participants.

General procedure

The general procedure was similar to experiment 4, except for the following: here the keys used to select an action were the 'right-arrow' and 'left-arrow' keys of a standard

keyboard, using the index and middle fingers of the right hand respectively. Participants reported the time by typing on the keyboard with their left hand. No visual feedback was presented following timing judgements as participant reports from experiments 4 and 5 indicated that they did not particularly attend to the visual feedback. Because it was redundant with the tone frequency, it was omitted in experiment 6.

This experiment only investigated action binding, and not tone binding, as action binding has been linked to outcome prediction mechanisms (Engbert & Wohlschläger, 2007) and to experience-dependent plasticity (Moore & Haggard, 2008). Further, excluding tone binding allowed us to increase the trial numbers in agency blocks without making the experiment excessively long.

Agency conditions

Besides the usual baseline measures, participants completed 5 blocks of 30 trials in the learning condition, and 5 in the chance condition, in pseudo-randomized order. As in experiment 4, in the learning condition one key delivered rewarded high tones with a probability of 0.8 and the other key with probability of 0.2. The high tone was always the 'correct' tone, and participants were told to learn which key was most frequently associated with the high tone. Participants were explicitly informed that reversals of the action-tone mapping would occur occasionally and unpredictably. These explicit instructions aimed to reduce the high inter-individual variability in performance found in experiment 4, by clarifying the task for poorer performers. Further, reversals now occurred after a variable number of trials (randomly either 6, 10 or 14 trials) so participants could not predict when they would occur. Run length was adjusted after the last reversal in the block to ensure the same number of trials for each participant. At the end of each block of the learning condition, if participants achieved a threshold of at

least 20 rewarded trials, they gained a bonus of 50p. A large blockwise reward was used in place of smaller trialwise rewards to avoid satiety after several successful trials, and to maintain motivation throughout.

In the random condition, the probability of hearing a high tone or a low tone was unrelated to the key chosen (50%/50%). Participants were explicitly told that their choice of action would not influence the tones they would hear. In the learning condition they were instructed to "find the good key, maximizing the number of high tones" while in the random condition they were told, "whichever action is chosen, it will have no influence on the following tone". Since learning could not be used to maximize reward in this condition, the number of high-tone trials did not lead to a monetary bonus. This arrangement ensured that participants were not incentivised to search for contingencies that did not exist. Although this creates a motivational difference between the two conditions, this bias is intrinsic to any reinforcement learning experiment (O'Doherty, 2014). Furthermore, at the beginning of each block, participants were explicitly told which condition they were in.

As before, participants reported the timing of their action. To further improve the precision of the measure, participants were instructed that at the end of each block they would receive an additional 25p if they improved the precision of timing estimates relative to the previous block. The standard deviation of their judgement errors was used to measure precision – note that this measure is independent of the mean timing judgement, and thus independent of action binding estimates. Therefore, in the learning condition, participants were rewarded for precision of timing judgements and for choosing the 'correct' key. In the chance condition, they were rewarded only for precision of timing judgements.

Baseline measures

Baseline conditions were the same as the action baseline conditions of experiment 4. Participants performed 2 baseline blocks of 20 trials each, at the beginning and end of the agency session. In baseline blocks, participants freely chose which of the two keys to press. No significant differences were observed in the perceived times of key presses in milliseconds for left and right hand responses ($F_{(1,29)}$ = 1.01, p = .319, η_p^2 = .018) or for pre- or post-experiment blocks measures ($F_{(1,29)}$ = .129, p = .721, η_p^2 = .002). Consequently, all baseline blocks were collapsed in further analysis.

Analysis

Action binding was calculated for each participant and each condition by subtracting the relevant mean baseline error from the error in agency trials. Agency trials were categorised according to three design factors: 1. whether the outcome on a given trial was a high or low frequency tone (associated with a rewarded or non-rewarded outcome respectively in the learning condition) 2. whether the trial was in the learning or random condition and 3. whether the outcome on the *previous* trial was a high or low frequency tone. Action binding data were then subjected to a 2x2x2 ANOVA.

4.4.3. Results

Performance

In the learning condition, participants demonstrated an ability to learn the correct action. As in experiment 4, the trial number after reversal had a significant effect on participants' proportion of correct choice ($F_{(5,145)}$ = 57.14, p < .001 , η_p^2 = .2). They

quickly returned to initial performance levels after a reversal event (see figure 4-6 in section 4.5.3. below).

Action binding

Action binding data are shown in figure 4-5. A 2x2x2 ANOVA revealed a highly significant main effect of condition, with stronger action binding in the learning condition compared to the chance condition ($F_{(1,29)}$ = 17.48, p < .001, η_p^2 = .376). There was no effect of current trial outcome ($F_{(1,29)}$ = .02, p = .896, η_p^2 = .001). Importantly, there was a significant main effect of previous trial outcome ($F_{(1,29)}$ = 14.56, p < .001, η_p^2 = .334) and also a highly significant interaction between learning condition and previous trial outcome ($F_{(1,29)}$ = 9.71, p = .004, η_p^2 = .251).

Simple-effect t-tests were used to further investigate this interaction. In the learning condition, non-rewarded outcomes significantly increased the action binding on the following trial compared to rewarded outcomes (simple-effect paired t-test: $t_{(29)}$ = 3.73, p < .001, *Cohen's d* = .685). This difference was numerically almost abolished, and became statistically non-significant, in the chance condition ($t_{(29)}$ = .46, p = .65; see figure 4-5).

No other interactions were significant (current trial outcome x condition: $F_{(1,29)}$ = .33, p = .57, η_p^2 = .011; current trial outcome x previous trial outcome: $F_{(1,29)}$ = 1.01, p = .32, η_p^2 = .034; current trial outcome x condition x previous trial outcome: $F_{(1,29)}$ = 0.13, p = .72, η_p^2 = .005).

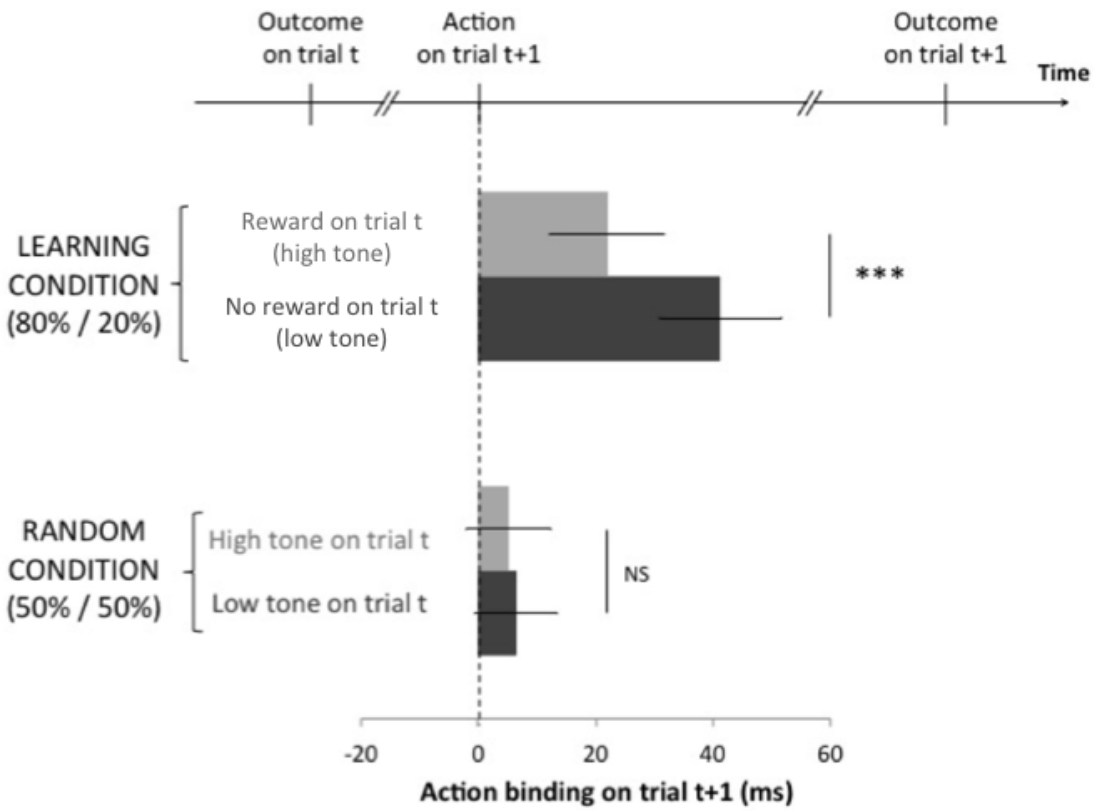

Figure 4-5. Mean action binding in ms following a rewarded (light grey) or non-rewarded (dark grey) outcome on the previous trial, for both chance and learning conditions. Error bars represent standard errors. Note that the high/low tones were associated with rewarded/non-rewarded outcome in the learning condition, but not in the chance condition. *** $p < 0.001$. Reproduced from Di Costa, Théro, Chambon, & Haggard (2017)

4.4.4. Discussion

With some changes in implementation, the post-error boost in action binding was replicated in the learning condition. Crucially, this effect was specific to this condition, and absent when participants could not learn stable action-outcome relations. These results therefore provide strong evidence that action binding reflects the ability to influence events through learning to improve one's own action choices. Critically, this learning depends on previous error feedback.

To further investigate this effect, a formal reinforcement learning model was used to investigate how the post-error boost in action binding is related to how people learn to maximize rewards.

4.5. Statistical Modelling of results from experiments 4 and 6.

4.5.1. Introduction

Consistent effects of learning were only obtained in experiments 1 and 3, thus only data from these experiments was modelled in this section.

Reinforcement learning models distinguish between the learning opportunities offered by errors and by rewards respectively. Interestingly, these two elements of learning are differentially expressed across the population. Negative learners are better at avoiding negative outcomes while positive learners are better at choosing positive outcomes. Interestingly, the EEG feedback-related negativity (FRN) evoked by an error signal was found to be larger in negative learners than positive learners (Frank, Woroch, & Curran, 2005). In the following analysis, it was predicted that the post-error boost in action binding might be positively correlated with participants' bias to learn more from negative than from positive outcomes.

4.5.2. Method and model

An established model of reinforcement learning was fitted to investigate whether inter-individual variance in asymmetric learning is correlated to the post-error boost in action binding. According to the reinforcement-learning algorithm, each of the possible actions (choosing the left or right button) was associated with an internal value called

an action value (Sutton & Barto, 1991). The values themselves are hidden, but are thought to drive choices between alternative actions.

Value updating

The model is based on the concept of prediction error, which measures the discrepancy between actual outcome value and the expected outcome for the chosen action (i.e. the chosen action value):

$$\delta(t) = \text{Outcome}(t) - \text{Value}_{\text{chosen}}(t)$$

Prediction error is then used to update the value of the chosen action. The values were set as 0.5 at the beginning of each block. As this analysis was concerned with the specific effect of rewarded and non-rewarded outcomes, two different learning rates, α+ and α-, were calculated to reflect different updating processes after a positive or negative prediction error (Lefebvre et al., 2016, Niv, Edlund, Dayan, & O'Doherty 2012). This asymmetrical model therefore accounts for individual differences in the way participants learn from positive and negative outcomes:

$$\text{Value}_{\text{chosen}}(t+1) = \text{Value}_{\text{chosen}}(t) + \begin{cases} \alpha^{+} \times \delta(t) & \text{if } \delta(t) > 0 \\ \alpha^{-} \times \delta(t) & \text{if } \delta(t) < 0 \end{cases}$$

The action values of the two possible actions were normalized by keeping their sum constant.

A reduced model was also constructed with only one learning rate for both rewarded and non-rewarded outcomes. The Bayesian Information Criterion (BIC) comparison showed that the BIC of the two learning rate model was significantly lower than the one rate learning model in experiment 4 (paired t-test : $t_{(15)}$ = 2.98, p = .01) but not experiment 6 (paired t-test : $t_{(29)}$ = 1.003, p = .32). However, the Aikake Information

Criterion (AIC) comparison showed that the AIC of the two learning rate model was significantly lower than the AIC of the one learning rate model for both experiment 4 (paired t-test : $t_{(14)} = 4.56$, $p < .001$) and experiment 6 ($t_{(29)} = 2.37$, $p = .025$). The model with two learning rates (α+ and α-) was thus considered the best fitting model.

In experiment 4, the chosen model (two learning rates) better predicted the data 10 participants, while the alternative model (with one learning rate) better predicted the data of 5 participants. In experiment 6, the chosen model outperformed the alternative for all 30 participants.

Decision rule

In the model, the action with the higher action value is more likely to be selected. The probability to choose an action will depend on the two action values and on the "inverse temperature" parameter β, which represents the strength of the action values' effect on action selection:

$$P_{\text{choosing Left}} = \frac{e^{\beta \times \text{Value}_{\text{Left}}}}{e^{\beta \times \text{Value}_{\text{Left}}} + e^{\beta \times \text{Value}_{\text{Right}}}}$$

Parameter fitting and simulations

The model parameters were fitted based on participants' choices on each trial. The three parameters fitted were: the two learning rates, α+ and α-, and the inverse temperature β. They were fitted independently for each participant, on the data from the learning condition in experiments 4 and 6. The best parameters chosen were those that maximized log likelihood (LLH), defined as the sum of the log of the model's fit to participant's action choices. Thus, LLH values close to 0 indicate a good model fit. A splice sampling procedure was used to test the different possible combinations of

parameters (Bishop, 2006). More precisely, using three different starting points drawn from uniform distributions for each parameter, 10,000 iterations of a gradient ascent algorithm were performed to converge on the set of three parameters that best fitted the data.

4.5.3. Modelling results and discussion

In experiment 4, the mean learning rates (α+ and α-) across participants were, respectively, .89 and .48 and the mean inverse temperature (β) was 3.65. In experiment 6, the mean learning rates (α+ and α-) across participants were, respectively, .67 and .51 and the mean inverse temperature (β) was 4.71.

From the fitted parameters, the simulated model's choices were a generally good match with participants' behaviour, see figure 4-6. The mean probability of the model selecting the same action as the participant was as follows: experiment 4: M = .73, SD = .07; experiment 6: M = .76, SD = .09). Thus the reinforcement-learning model seemed to accurately reflect participants' learning processes. Similarly to Lefebvre et al. (2016), overall higher learning rates were found for rewarded outcomes than for non-rewarded outcomes (experiment 4: α+ M = .89, SD = .13, α- M = .48, SD = .14, paired t-tests, $t_{(14)}$ = 9.15, p < .001; experiment 6: α+ M = .67, SD = .27, α- M = .51, SD = .23, $t_{(29)}$ = 3.26, p = .003), justifying the use of an asymmetrical model.

The normalized learning rate asymmetry (Lefebvre et al., 2016; Niv et al., 2012), defined as $(\alpha^- - \alpha^+)/(\alpha^- + \alpha^+)$ was calculated to investigate whether the post-error agency boost could be related to the outcome-specific learning-rate. The post-error boost in

action binding was defined as the difference between action binding after a non-rewarded outcome and action binding after a rewarded outcome, as before. For experiment 4, there was no relation between post-error agency boost and normalized learning rate asymmetry ($t_{(13)}$ = -.66, p = .518, R^2 = .03). However, there was a positive correlation between post-error agency boost and normalized learning rate asymmetry in the learning condition of experiment 6, ($t_{(28)}$ = 5.6, p = .026, R^2 = .17, see figure 4-6), implying that individuals who learn from errors also show a strong post-error agency boost. The absence of any effect in experiment 4 may reflect the lower statistical power, and may also reflect the very restricted inter-individual variability in learning rate asymmetry (asymmetry in experiment 4: M = -.31, SD = .14; in experiment 6: M = -.15, SD = .32).

Finally, it was considered whether other confounding factors, in addition to normalized learning rate asymmetry, could predict individual variability in post-error agency boost in experiment 6. In particular, one alternative view hypothesizes that the post-error agency boost could merely reflect saliency of rare error events, akin to the non-specific alerting effect of an oddball, rather than any relation between errors and learning. This alternative model also predicts a negative relation between an individual's post-error agency boost and the frequency of their errors, yet no such relation was found ($t_{(28)}$ = .53, p = .603, R^2 < .001), and the sign was not as predicted. The general influence of contingency is further explored in the following chapter.

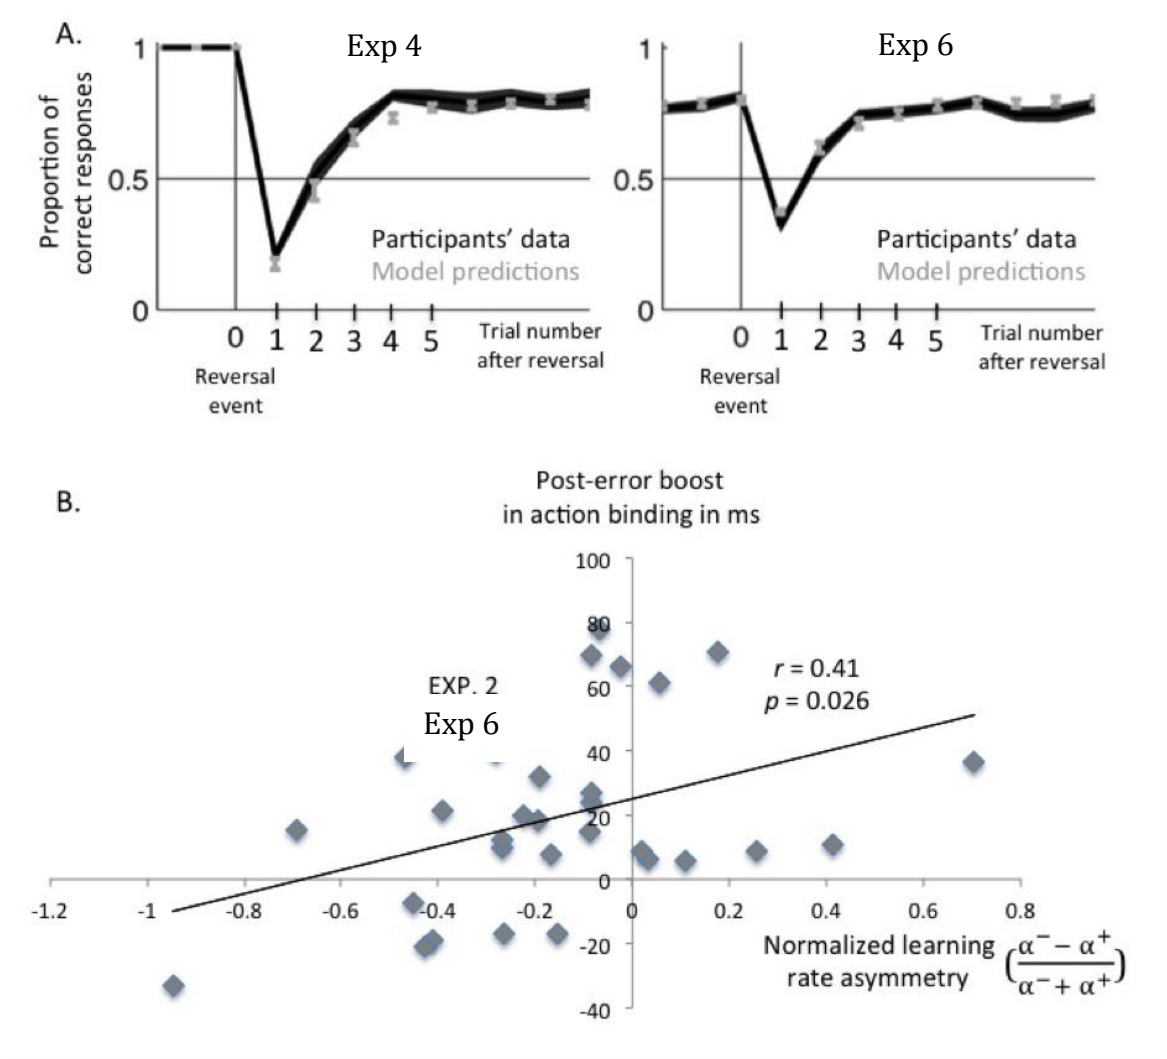

Figure 4-6: Proportion of correct responses before and after a reversal event for experiments 4 (left panel) and 6 (right panel). Participants' data are in black and predictions of the reinforcement learning model are in grey. B. Post-error boost in action binding plotted against the normalized learning rate asymmetry for Experiment 6. Reproduced from Di Costa, Théro, Chambon, & Haggard (2017)

4.6. General Discussion

4.6.1. Chapter overview

Experiments in this chapter examined two measures of intentional binding: the binding of actions toward their consequences and the binding of outcomes/tones back towards the events that caused them, in both learning and non-learning contexts.

4.6.2. Outcome/Tone binding

With regard to outcome binding, experiment 6 replicated the difference in outcome binding found in experiment 4. Outcome binding was stronger in non-rewarded trials compared to rewarded trials. This effect was small and contrary to previous results (Yoshie & Haggard, 2013; Takahata et al., 2012) so its meaning remains unclear. Those studies suggested that the well-known self-serving bias (Bandura, 1989) might influence not only explicit attributions of agency, but also implicit measures of the basic experience of agency. However, the studies presented in this chapter add an additional, important element of learning, which those earlier studies lacked. The effects of learning from errors appear to replace or outweigh the effects of valence. In the novel paradigm presented here, errors provided important evidence for learning what action to perform next. In contrast, the valence of outcomes in previous experiments was completely predictable, and unrelated to action choices. The observed effects of valence on outcome binding were not replicated in any further experiments in this thesis and are not discussed further.

4.6.3. Action binding

In all three experiments presented in this chapter, actions were more strongly bound towards their outcomes following a non-rewarded trial. This effect was observed wherever learning could be used to maximise outcomes. In the non-learning condition of experiment 6, participants were explicitly told that they would not be able to maximise rewards through learning and no post-error effect was observed.

One may therefore consider that the post-error boost to the sense of agency is indeed a feature of learning. Information provided by prediction errors is only useful if another

attempt will be made to achieve the same goal. It therefore follows that the effect of errors on the sense of agency would only be observed in a learning context, as demonstrated in experiment 6. The *reliability* of prediction errors, as manipulated in experiment 5, might inform the selection of actions, but not necessarily alter the feeling of control over the actions themselves.

4.6.4. The post-error agency boost – potential explanations

The concept of 'cognitive control' refers to the control and monitoring of cognitive resources to achieve successful task performance. Errors signal a failure of effective control, and trigger a number of adaptations, notably 'post-error slowing' (Danielmeier & Ullsperger, 2011). Post-error slowing is classically associated with increased caution in action selection after errors (Dutilh et al., 2012). The relation between the post-error agency boost and post-error slowing remains unclear. However, it seems unlikely that a mere transient increase in availability of general cognitive resources devoted to action selection, as suggested by conflict adaptation theories, can explain the increase in post-error action binding. A general boost in attention following an error would be expected to cause a general increase in precision of timing judgements, reducing judgement errors, and therefore reducing both action binding and tone binding effects – yet a specific increase in judgement errors was observed for actions only. Instead, the post-error boost may reflect a specific strategic adaptation to the information value of the trial following an error. This adaptation reflects the fact that errors may be highly informative for future actions. Strongly linking actions to outcomes on the trial following an error may be important for guiding future action choices.

4.6.5. Potential confound of frequency

One may argue that the post-error agency boost could merely reflect saliency of rare error events, akin to the non-specific alerting effect of an oddball, rather than any relation between errors and learning. This alternative model predicts a relation between an individual's post-error agency boost and the frequency of errors. No such effect was observed in the experiments presented here. Nevertheless this potential confound, and related questions, are explored more thoroughly in chapter 5.

4.7. The post-error agency boost – preliminary conclusions

Taken overall, cognitive control mechanisms engaged when people make errors may have two distinct effects: an increase in cognitive resources to restore performance, and an increase in the experiential link between action and outcome. The latter effect could trigger a post-error boost in agency. However, the studies presented here cannot identify for certain the direction of any causal relation between the post-error agency boost and learning from errors.

Chapter 5: Influence of frequency and predictability of outcomes in the novel paradigm

5.1. Introduction

This chapter will investigate two questions that arise from the preliminary work described in chapter 4.

Experiment 7 examined the potential influence of outcome frequency as a potential confound. In the preceding experiments, rewarded trials appeared more frequently than non-rewarded trials, as participants learned to produce the more-rewarded action. Stronger binding for negative compared to positive outcomes could potentially be due to familiarity, rather than reward value. Experiment 7 was designed to control for this alternative possibility.

Experiment 8 investigated the nature of the mechanisms underlying the binding effects observed thus far. Time judgments in intentional binding tasks are typically given after the participant has been made aware of the outcome. It is therefore possible that the nature (or mere presence) of the outcome influences participants' perception retrospectively. This is a view favoured by many researchers (see for example Wegner & Wheatley, 1999). However, a *prospective* mechanism would have an influence on binding that is insensitive to the nature of the outcome. Indeed, any effects that are associated with this prospective mechanism should persist even in the absence of an anticipated outcome. This hypothesis is tested in experiment 8 (section 5.3. below).

5.2. Experiment 7. Testing the frequency confound hypothesis

5.2.1 Introduction

This experiment was designed to examine the potentially confounding effect of outcome familiarity in the novel paradigm described in chapter 4. Dummy trials were inserted into the trial schedule and matched negative trials in frequency but carried no information relevant to the task. It was predicted that action binding on trials following negative outcomes would be stronger than on dummy trials. It was also predicted that action binding on trials following negative outcomes would be stronger than on trials following positive outcomes, replicating the effects observed in experiment 4, 5 and 6.

5.2.2. Method

Participants

Sixteen participants (9 female, all right-handed, mean age = 24 years, age range = 19-43 years) completed the experiment and were paid £7.5/hour. All participants gave informed written consent prior to commencing the experiment. All participants reported normal or corrected to normal vision and hearing.

Apparatus, design and analysis

The experiment was based on experiment 4, however dummy trials were randomly inserted into the trial schedule. A dummy trial always produced a neutral tone, intermediate in frequency between those signalling reward (2000 Hz) or penalty (500 Hz). Participants were instructed that dummy trials appeared randomly, and did not influence the probability of outcomes in following trials. The appearance of dummy trials was yoked to the occurrence of a penalised outcome but did not necessarily occur

immediately after. Penalised and dummy trials therefore occurred with precisely the same frequency within a block.

To balance the valence of outcomes, successful trials were rewarded (10p) while unsuccessful trials incurred a penalty (-10p). Dummy trials were neither penalised nor rewarded.

As in experiment 6, reversal events occurred automatically, every 6-8 trials. Participants completed six blocks of fifty trials for a total of 300 trials.

Analysis of baseline trials

No significant differences were observed in the perceived times of key presses in milliseconds for left and right hand responses, forced- or free-choice, or for pre- or post-experiment blocks measures ($p > .05$ for all comparisons). Consequently, all action baseline blocks were collapsed in further analysis.

Analysis of operant trials

Perceptual shifts were calculated as in experiment 4.

5.2.3. Results

Simple paired t-tests were used to analyse the main effect of the previous outcome on action binding. Action binding following a penalty was significantly higher than action binding following a reward ($t_{(15)} = 2.6$, $p = .02$, *Cohen's d* = .64, see figure 5-1). Crucially, the difference between action binding following penalty and dummy outcomes was marginally significant ($t_{(15)} = 1.8$, $p = .09$, $d = .45$, see figure 5-1).

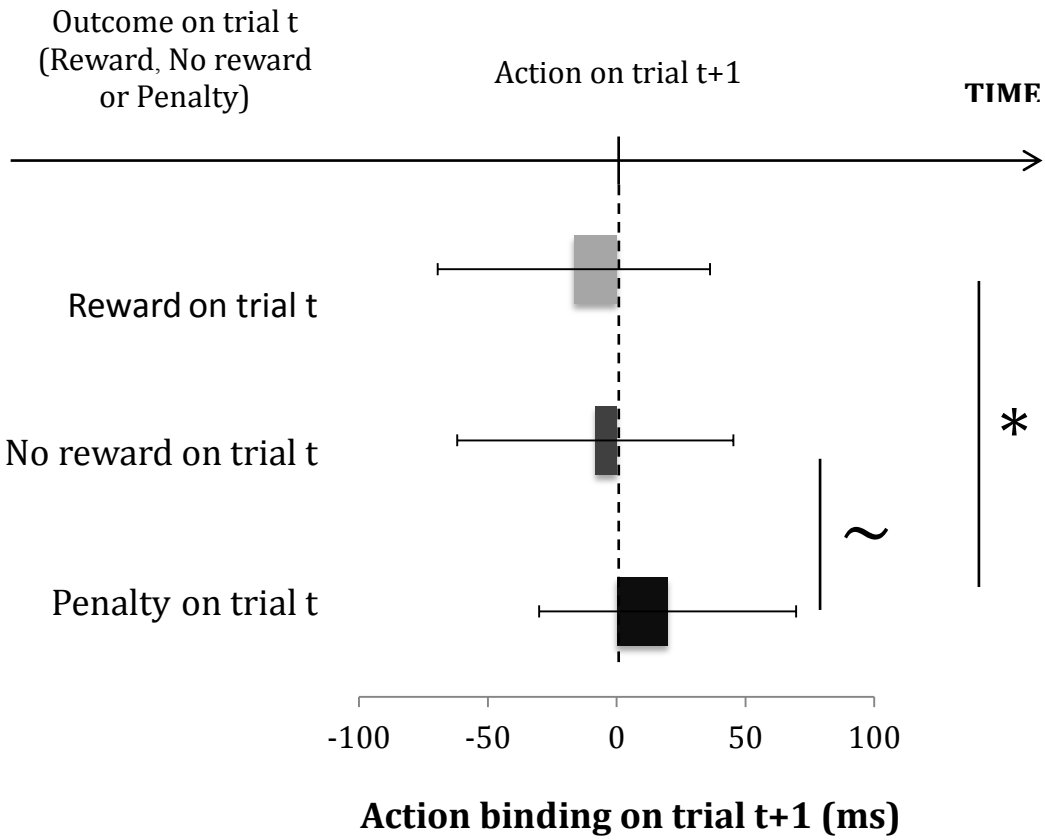

Figure 5-1. Action binding in experiment 7.

5.2.4. Discussion

This experiment tested the possibility that stronger action binding following a negative reward might be driven by differences in frequency of exposure between more and less rewarding outcomes.

Action binding following a penalty was stronger than following a rewarded outcome, replicating the post-error agency boost observed in Chapter 4. More importantly, although penalised outcomes and non-rewarded dummy outcomes occurred with the same frequency, penalised outcomes tended to produce more subsequent action binding than non-rewarded outcomes. Thus the effect of previous outcome valence on current action binding is not convincingly explained by an "oddball" effect.

5.3. Experiment 8: Testing for prospective mechanisms of the sense of agency.

5.3.1. Introduction

As discussed in the General Introduction to this thesis (section 1.4.5.), research has suggested that the mechanisms underpinning the sense of agency may be either prospective or retrospective in nature. Implicit measures such as intentional binding may therefore capture aspects of the sense of agency that are associated with signals preceding the goal-directed action, or signals that arise once the outcome of the action is known, or both.

The post-error boost to action binding observed in four experiments thus far is a sequential effect. The novel paradigm described in this thesis therefore offers a new opportunity to investigate potential *prospective* influences on the sense of agency. If the post-error action boost is observed even in the absence of an anticipated outcome, one may conclude that prospective mechanisms cannot be ruled out as an explanation for the post-error action boost. Experiment 8 was designed to test this hypothesis, by including trials where the anticipated outcome does not occur.

5.3.2. Method

Participants

Eighteen participants (13 female, all right-handed, mean age = 25 years, age range = 21-51 years) completed the experiment and were paid £7.5/hour. All participants gave informed written consent prior to commencing the experiment. All participants reported normal or corrected to normal vision and hearing.

Apparatus, design and analysis

This experiment was based on experiment 7, in which dummy trials were randomly inserted into the trial schedule, occurring with precisely the same frequency as penalised trials within each block.

Furthermore, we varied the reliability of the outcome. On one third of all trials, the tone did not occur. Participants were simply asked to judge the time of their action, and the experiment continued.

Note that in this experiment dummy outcomes are distinct from absent outcomes (see figure 5-2).

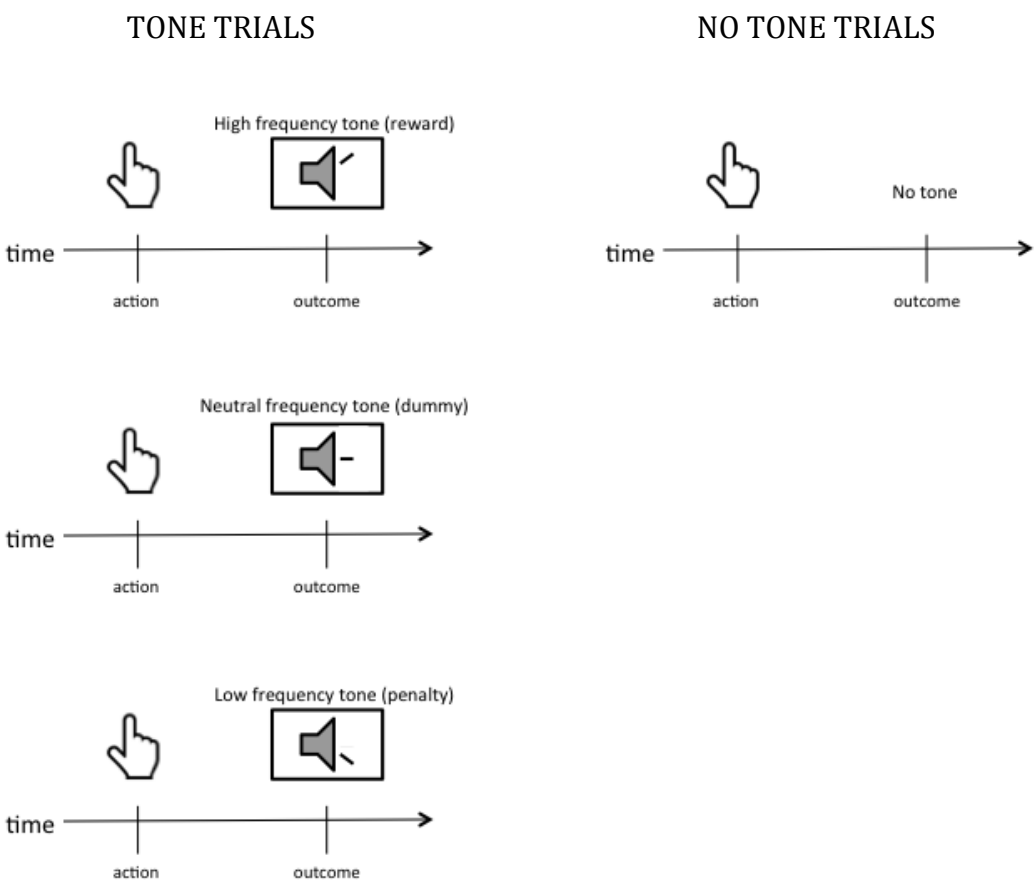

Figure 5-2: Trial types in experiment 8. Left panel: on two thirds of trials, a tone is heard following the action event, according to a probability schedule Right panel: on one third of trials, no tone is heard following the action event. Note the distinction between a dummy trial (left, center) and a no-tone trial (right, top).

Analysis of baseline trials

No significant differences were observed in the perceived times of key presses in milliseconds for left and right hand responses, forced- or free-choice, or for pre- or post-experiment blocks measures (p > .05 for all comparisons). Consequently, all action baseline blocks were collapsed in further analysis.

Analysis of operant trials

Action binding was significantly stronger following a negative (penalised) outcome than following a dummy outcome ($t_{(17)}$ = 2.15, p = .047), suggesting that frequency was not a confounding factor in action binding results and replicating the trend observed in experiment 7. All trials following a dummy outcome were therefore excluded.

Furthermore, all trials following a trial with no tone were excluded.

3.3.3. Results

Action binding results were submitted to a 2x2 repeated measures ANOVA with factors: previous trial feedback (rewarded/penalised) and current trial type (tone present/absent). Action binding was significantly stronger following a penalised trial than a rewarded trial ($F_{(1,17)}$ = 6.15, p = .02, η_p^2 = .266). There was no effect of trial type on action binding ($F_{(1,17)}$ = .21, p = .65) but there was a significant interaction ($F_{(17)}$ = 5.94, p = .026, η_p^2 = .259). Post-hoc t-tests confirmed significantly more action binding after negative outcomes than after positive outcome in tone absent trials ($t_{(17)}$ = 2.99, p = .008) but not in tone present trials ($t_{(17)}$ = .06, p = .95), see figure 5-3.

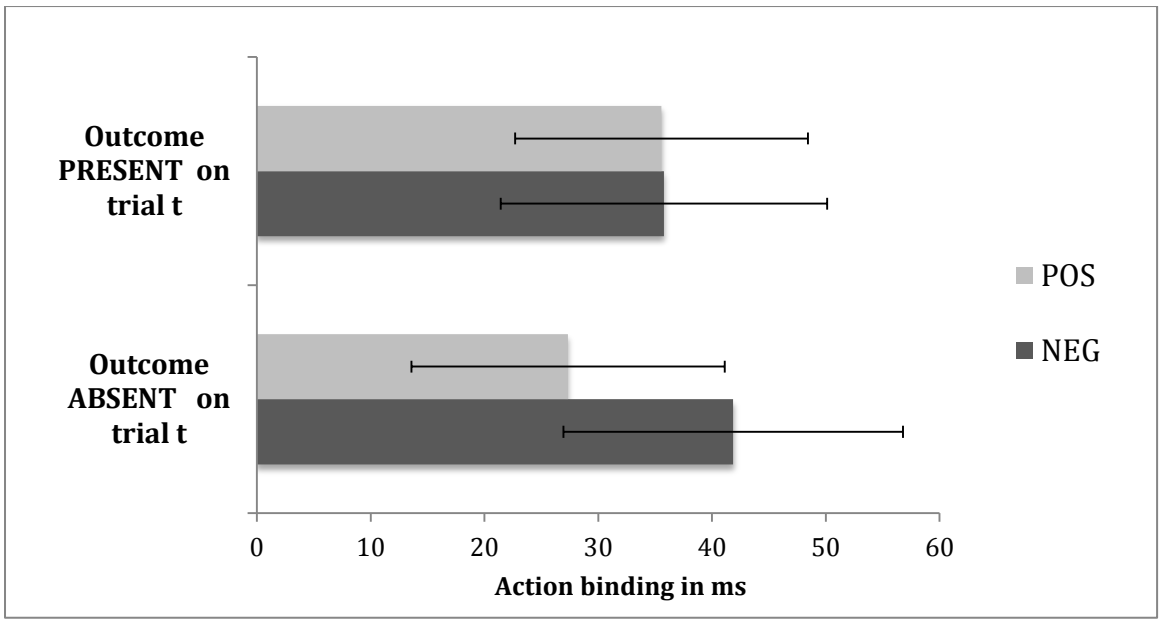

Figure 5-3. Action binding in experiment 8.

3.3.4. Discussion

Experiment 8 was designed to investigate the mechanisms thought to underlie SoA. Specifically, anticipated but absent outcomes were used to probe the prospective/retrospective aspect of SoA, as measured with Intentional Binding.

The presence or absence of outcomes did not have a significant effect on action binding results. There is therefore no evidence that the mechanisms under investigation are entirely retrospective.

Indeed, the post-error boost to action binding described in previous experiments (and replicated here) was only detectable when the tone was absent. This was a surprising result, as the effect was expected to persist regardless of outcome presence or absence. The occurrence of outcomes in an unreliable environment may take on a greater associative importance, as the brain attempts to gather information while it is available. The salience of such outcomes may then efface the post-error agency boost, as seen here.

Results in this experiment suggest that the mechanisms giving rise to this effect are at least partly *prospective*. Although there is not yet a comprehensive understanding of the precise causes that give rise to agency, one can describe certain features of the implicated processes. The results obtained here rule out an entirely retrospective explanations. Furthermore, the use of the novel paradigm used here highlights an important environmental feature (outcome obtained) that can prospectively influence the subjective experience of future associated actions.

5.4. General Discussion

5.4.1. Results obtained in this chapter

The previous chapter introduced a novel paradigm to study SoA. An intentional binding task was embedded in a reversal-learning environment. For the first time, sequential effects on binding were identified and described. The two experiments described in this chapter were proposed to investigate the observed boost to action binding following a negative outcome. Experiment 7 was designed to control for the potentially confounding effect of outcome frequency. Experiment 8 made use of absent outcomes to delineate retrospective and prospective aspects to the sense of agency, including the effect described above.

The relative infrequency of negative outcomes did not appear to be a convincing explanation for the post-error action boost. In both experiments, negative outcomes tended to produce more action binding than dummy outcomes, despite occurring with precisely the same frequency.

Furthermore, the post-error action boost arises from mechanisms that cannot be entirely retrospective in nature. In experiment 8, the effect was observed on trials where outcomes were anticipated but did not occur.

3.4.2. Synthesis of results from Chapters 4 and 5.

The capacity to choose between actions in order to obtain rewards seems essential to functional behaviour outside the laboratory, and the sense of agency would be an important experiential component of that capacity. The paradigm described here forced participants to continuously learn relations between actions and outcomes. Previous studies showed that intentional binding is sensitive to economic (Takahata et al., 2012) and affective (Yoshie & Haggard, 2013) valence, but these studies did not address learning and decision-making aspects. The paradigm introduced in chapter 4 has been used to demonstrate for the first time how outcome success or failure influences the sense of agency in a dynamic learning environment.

While sense of agency is usually defined as the feeling of controlling one's actions and their consequences (Haggard & Chambon, 2012), few studies have investigated the discriminative ability to control outcomes. One suggested that action-outcome relations had no effect on intentional binding (Desantis, Hughes, & Waszak, 2012). Unlike previous studies, the experiments presented here involved an element of reward-based decision-making. Thus, it may be argued that action binding is a valid implicit measure of *purposive* outcome agency, i.e. the ability to generate one particular external event, rather than another, through one's own motivated, endogenous action.

People normally make actions for a reason. That is, they choose actions to achieve a desired outcome. They then monitor and evaluate whether the action succeeded or

failed in achieving the outcome. Thus, one might intuitively expect a link between adaptive behaviour and sense of agency, yet these two traditions in action control have evolved through largely separate research literatures. The data presented here suggest that when people experience negative outcomes, they feel more control, not less, in the next trial. This may initially seem counterintuitive, but it is strongly consistent with the view that sense of agency is related to acquiring and maintaining control over external events.

3.4.3. Conclusions regarding the post-error action boost

SoA has an important functional role in adaptive behaviour. Negative feedback might increase participants' feeling of agency for a short time, because action failures should strongly motivate the requirement to act appropriately in the future. SoA could be understood in the context of motivation to improve performance on subsequent actions.

The reliability and general effect size of the post-error action boost will be subjected to meta-analysis, described in the following chapter.

Chapter 6: Reliability and effect size of the post-error boost to sense of agency – a meta-analysis.

6.1 Introduction

In the preceding two chapters, intentional binding was embedded in a reversal-learning task. It was observed that action binding was stronger on trials following an undesirable outcome, such as a monetary loss. It was speculated that this novel, sequential effect may be a marker of cognitive processes such as learning, error monitoring and action selection.

In chapter 5, the effect proved robust to replication. It was also shown that the relative infrequency of negative outcomes was not a convincing explanation for the observed boost. However, it was also noted that the effect is quite unlike results from previous experiments using intentional binding and valenced outcomes (Yoshie & Haggard, 2013, Takahata et al., 2012).

It is therefore prudent to consider the reliability and relative strength of the post-error boost using a formalised method of inter-study comparison. This chapter will examine the combined results of the experiments described in chapters 4 and 5, and other studies not described in full in this thesis, using meta-analysis.

As a statistical tool, meta-analysis is more frequently employed for clinical trials than for laboratory experiments such as those described in this thesis. As a result, meta-analysis is commonly used for between-subject designs more than for within-subject designs. This is an important consideration in the present work, as the effect size for a within-subject design can be calculated from the difference between the two conditions

of interest, whereas the effect size for between-subject designs is based on estimates of variability in each of the two conditions. Meta-analyses of within-subject studies typically report effect sizes based on between-subject effect sizes (Lakens, 2013). This is partly motivated by a desire for easy comparability with other meta-analyses but also because complete variability data is rarely provided in published reports of within-subjects designs. Nevertheless, employing such methods may result in poorer estimates of effect size.

Relevant to the present analysis, if the variability of a within-subjects experiment arises from factors common to all conditions, then the variability of difference scores may be substantially lower than the combined variance of individual conditions. A method of calculating error variance for between subject designs would therefore underestimate the within-subject effect size. Additionally, and unusually for a meta-analysis, all data from each experiment included here was available. Accordingly, the analysis described in this chapter will employ a measure of effect size recommended for within-subjects designs. Specifically, variability of difference scores is calculated using the average standard deviation of all conditions (Lakens, 2013).

6.2. Method

6.2.1. Study Selection for Multi-study analysis

There are few studies that have been conducted using the same paradigm described in the present work. Previous studies using intentional binding could not be included due to the absence of any sequential dependency between trials. Similarly, previous reversal learning experiments have not employed intentional binding as a measure of agency and are not therefore comparable.

To be considered relevant, studies must have measured action binding as a dependant variable. Furthermore, it must have been possible for participants to learn from the outcome of an individual trial to improve performance on subsequent trials. Data from patient studies were excluded, as were experimental trials from brain stimulation experiments.

All relevant studies from this thesis are included. A further experiment was included that was conducted by colleagues independently and in collaboration with the author.

6.2.2. Description of studies

Experiments from this thesis are described in full in their respective chapters. Experiments 4, 5 and 6 are described in chapter 4. Experiments 7 and 8 are described in chapter 6. The following experiment is not described in this thesis:

Di Costa, Khalighinejad & Haggard (2016, unpublished): Design and execution as in experiment 7, however Transcranial Magnetic Stimulation (TMS) was applied between trials. Only sham condition data are included here.

6.2.3. Meta-analysis method

A fixed effects meta-analysis was conducted using the steps described by Lipsey and Wilson (2001). Paired-sample t-tests were performed for each study, comparing action binding on trials following positive and negative outcomes. Effect sizes were obtained using Cohen's d_{AV} and corrected for sample size biases using Hedges's G_{AV} (Lakens, 2013). These measures are easily compared with between-subject effect size measures that may be more familiar. Furthermore, other effect size measures may provide conservative results when correlations are high across conditions (Lakens, 2013).

The purpose of this analysis was to determine the strength and reliability of an observed trend across experiments. The direction of effect was therefore known: binding was expected to *increase* following a negative outcome on the previous trial. G_{av} results across all studies were therefore submitted to a one-tailed Z-score test to determine difference from zero.

The use of fixed- or random-effects analyses is an important statistical consideration (Hedges & Vevea, 1998). While fixed effects models assume that all experiments share a common effect size, random effects models assume that the effect size varies across experiments (Smith & Alloy, 2009). Previous meta-analyses of cognitive effects have argued in favour of fixed-effects models when the task and experimental environment were the same (see for example Horvath, Forte & Carter, 2015). A fixed-effects model was considered appropriate in the present analysis. All experiments were conducted in the same laboratory, using the same software to present stimuli. Although individual parameters varied by experiment, the fundamental design of the task remained the same. Lastly, subjects in each experiment were drawn from the same subject pool. It was therefore considered that reasonably homogenous results would be obtained. Cochran's Q was used to investigate this assumption, by testing heterogeneity between experiments (Higgins, Thompson, Deeks, & Altman, 2003).

6.3. Results

6.3.1. Descriptive data

The data used in the meta-analysis are summarised in table 6-1. For each experiment, action binding results were extracted and categorized according to the outcome on the previous trial. Rewarded trials were considered 'positive' and penalties (or non-

rewarded trials in experiments without penalties) were considered 'negative'. Trials following a dummy or absent outcome (in experiments that included those features) were excluded.

		Negative outcome on trial t-1		Positive outcome on trial t-1	
	N	M	SD	M	SD
Experiment 4	15	87.24	62.84	62.98	49.23
Experiment 5	16	50.38	52.11	44.25	48.58
Experiment 6	30	41.57	57.16	22.00	53.75
Experiment 7	16	19.73	199.61	-16.63	211.58
Experiment 8	18	32.06	94.17	22.79	107.57
Di Costa et al.	12	67.74	47.77	61.49	43.63

Table 6-1: Sample size (N), Mean (M) and Standard Deviation (SD) in ms, of relevant conditions in all included experiments.

6.3.2. Analysis of effect sizes

The general impression that emerges from individual study results largely supports the concept of a post-error boost to the sense of agency, as measured with intentional binding. Statistical comparisons are summarised in table 6-2. Four of the six studies in the analysis showed a significant increase in action binding following a negative outcome when conditions were directly contrasted.

	Observed effect	Statistical test	Effect size (G_{AV})
Experiment 4	Significant increase	$t_{(14)} = 3.4$, $p = .005$	.43
Experiment 5	Significant increase	$t_{(15)} = 2.28$, $p = .038$	.2
Experiment 6	Significant increase	$t_{(29)} = 3.73$, $p < .001$	.35
Experiment 7	Significant increase	$t_{(15)} = 2.6$, $p = .02$	.17
Experiment 8	No significant result	$t_{(17)} = 1.65$, $p = .117$	.09
Di Costa et al.	No significant result	$t_{(11)} = 1.16$, $p = .27$	.13

Table 6-2. Effect of negative outcome on action binding on the following trial in all included experiments.

The results obtained were then used to determine if there was statistical support for a general post-error boost to the sense of agency. Results were also used to investigate if

there was any heterogeneity between experiments beyond what could be expected from sampling error.

A forest plot of the included studies and an overall effect size is shown in figure 6-1. Note that unlike the paired-sample t-tests, the recommended statistic in the meta-analysis (G_{av}) is insensitive to the correlation between samples. Consequently, the plot shows confidence intervals that include zero (Lakens, 2013).

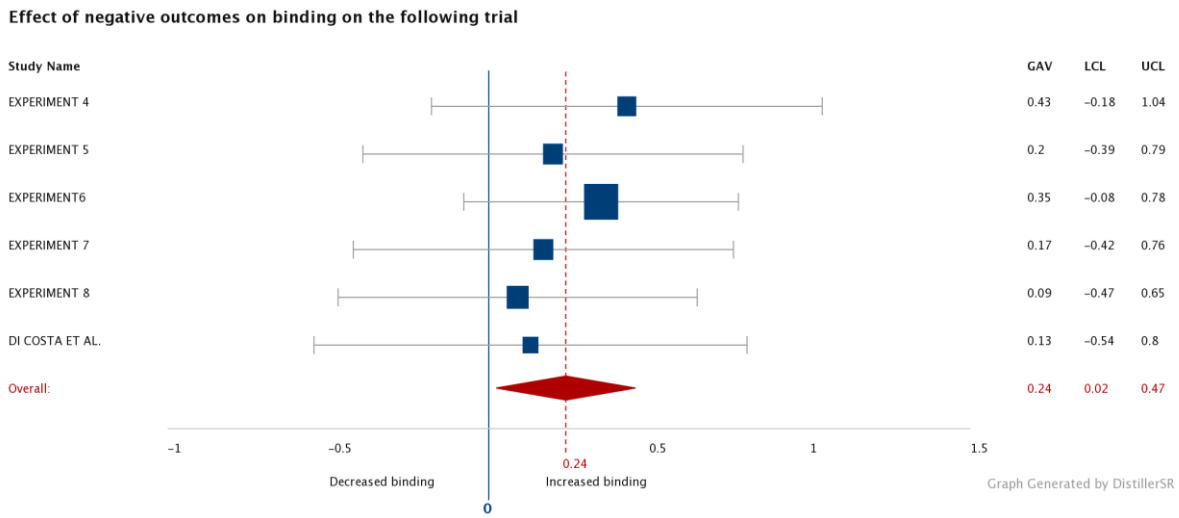

Figure 6-1. Forest plot of results from individual studies and overall effect size. Note that the recommended statistic (G_{av}) is insensitive to the correlation between samples. Consequently, the plot shows confidence intervals that include zero.

The overall effect of stronger action binding following a negative outcome was found to be significant: Mean (SE) G_{av} = 0.24 (0.14), z-test = 1.76, p = 0.039, 95% CI = [0.02 0.47].

There was no evidence for any heterogeneity in the results, beyond what would be expected by chance: $Q_{(5)}$ = .76, p > .1.

6.4. Discussion

This analysis compared the results of six studies, each measuring intentional (action) binding within a reversal learning paradigm. The effect of interest was an increase in action binding following a negative outcome. Effect sizes were calculated and compared across experiments to determine the strength of this sequential effect.

The overall impression is of a small but reliable effect in the predicted direction i.e. an increase in action binding following negative outcomes. The overall effect size was significantly higher than zero, indicating that when all relevant studies were considered, there was robust statistical support for the effect described. Furthermore, the effect size confidence intervals suggest the effect is noisy, but reliable. Intervals were relatively large, but unimodal, skewed above zero and with very few outliers.

Meta-analysis has its origins in the systematic appraisal of clinical trials. As such, the recommended methods do not necessarily transfer neatly to experimental settings such as those described in this work. In ideal circumstances, meta-analysis includes several large studies all using a common intervention with a precise measurable outcome. In contrast, the experimental psychology literature comprises many low-powered studies. Variations in equipment, attention, personality and experimenter technique may influence effect size.

The analysis of heterogeneity indicated that the results obtained were highly homogenous. This was unsurprising given that all included studies were conducted by or in collaboration with the author, using the same experimental software and generally adhering to the same experimental protocol. Furthermore, although the participant samples were balanced for gender, and diverse in terms of age and ethnicity, all

participants were recruited from the same subject pool. This degree of homogeneity was beneficial for confidence in inter-study comparison, but limits the generalizability of findings. This limitation is not comprehensively addressed in this thesis, however data from older participants and patients groups are presented and discussed in chapters 2, 3 and 7.

The meaning and significance of the post–error boost to action binding has been discussed in previous chapters. The purpose of this chapter was to provide evidence that the observed trend from several low-powered studies was robust to a systematic inter-study analysis. The results show in a quantitative manner that the effect is reliable and therefore suitable for making further experimental predictions.

It is unknown which cognitive and physiological mechanisms underpin the post-error boost. Experiment 6 provided evidence which suggests that it arises in contexts where monitoring outcomes is critical in learning what to do next. However it remains unclear whether the post-error agency boost is a marker of learning processes specifically and not simply a phenomenon that arises when one is engaged in goal-based decision-making.

The following chapter will describe results obtained from a PD patient group and age-matched healthy controls, using the novel paradigm described here.

Chapter 7: Sense of agency and learning in Parkinson's disease

7.1. Introduction

Learning to select actions based on outcomes is a core cognitive capacity. Impaired learning and impulse control are features of Parkinson's disease (PD) that often result in poor performance in goal-based decision-making tasks. The role of agency in this context is poorly understood.

Chapters 4, 5 and 6 describe a reliable post-error agency boost in young healthy volunteers following negative outcomes. In experiment 7, the novel paradigm described in chapter 4 was used once again. Participants were PD Patients on medication and healthy age-matched controls.

7.2. Experiment 9: Sense of agency and learning in Parkinson's disease

7.2.1. Introduction

General time perception deficits in PD are described in the introduction to this thesis (section 1.3.9.) and in chapters 2 and 3. Relevant to this experiment, it was shown in experiments 1, 2 and 3 that PD patients performed with similar precision in a Libet task to healthy controls. However, intentional binding is a measure in the *shift* in time perception from one condition to another.

Moore et al. (2010) found that patients OFF medication gave temporal judgements similar to healthy controls in a binding task. However, binding in patients was significantly increased while ON medication compared to OFF medication.

This experiment compared the performance and time judgements of patients ON medication to those of healthy age-matched controls. Stronger action binding was expected in patients compared to controls, as a result of being ON medication. It was also predicted that patients would show degraded performance in the decision-making task, due to generally impaired learning. Finally, it was predicted that the post-error agency boost would be reduced in patients relative to healthy controls, as this effect appears to be strongly related to learning.

7.2.2. Method

Participants

This study was approved by the UCL Research Ethics Committee. All participants agreed to participate in the study and signed a consent form. They all reported normal or corrected to normal vision and hearing, and were paid a bonus contingent on performance in the task.

Eighteen PD patients (6 female, all right-handed, mean age = 60 years, SD = 8.7) attending the movement disorders outpatient clinic at the University Hospital of Messina, Italy, agreed to participate and were included in the study.

Inclusion criteria were: a diagnosis of PD according to the UK Parkinson's Disease Society Brain Bank criteria; treatment and clinical condition stable for at least 4 weeks prior to the study.

Exclusion criteria were: any major concurrent neurological or psychiatric disorders (no exclusions); or a score < 24 on the Montreal Cognitive Assessment (MOCA) (one exclusion); language comprehension abilities inadequate to understand the instructions of the task (one exclusion). We therefore analysed data from 16 patients.

Seventeen healthy, age-matched (50-80 years old) volunteers in both Messina and in London, who responded to an advertisement were also included in the study and served as control group.

For all patients the following demographic and clinical data were collected: educational level, age at study entry, age at disease onset, disease duration, as well as levodopa and dopamine-agonist equivalent daily dose (LEDD) (Tomlinson et al., 2010). Degree of motor impairment was rated by means of the Unified Parkinson's Disease Rating Scale part 3 motor examination (UPDRS-III). All patient were tested in their ''best ON'' state (the best motor state at peak effect after taking their usual medication dose).

All patients underwent a neuropsychological assessment including: the Beck Depression Inventory (BDI), the Hamilton Anxiety Rating Scale (HAM-A) and the Apathy Evaluation Scale (AES) to rate, respectively, depression, anxiety and apathy. The Questionnaire for Impulsive-compulsive Behaviour Disorders in Parkinson's disease (QUIP-RS) (Weintraub et al., 2012) was employed to evaluate the severity of ICB.

The Montreal Cognitive Assessment (MOCA) was employed to assess overall cognitive function; digit span forward and backward for working memory; phonological word fluency and categorical word fluency for executive functions and language; Matrix Cancellation Features Target (MCFT) for selective visual attention.

All clinical instruments are included in Appendix I.

Procedure

Materials and general procedure were as in experiment 4.

Baseline conditions were as in experiment 4 and were repeated at the end of the experiment.

Operant conditions were based on experiment 4, however reversals occurred after a variable number of trials (randomly either 8, 10 or 12 trials), as in experiment 6.

Given specific predictions based on previous work, only action binding was recorded.

Analysis

Judgement errors in baseline blocks and binding in operant blocks were calculated as in experiments 4-8.

A measure of perseveration was also taken, defined as the frequency of trials in which the participant switched keys (relative to the previous trial). A tendency to perseverate would therefore reveal itself in low rates of switching across the experiment.

7.2.3. Results

Clinical data of the PD population are shown in table 7-1.

Patient ID	UPDRS-III	QUIP-RS	BDI	HAM-A	AES	MOCA
P_01	10	0	4	4	4	28
P_02	6	26	4	9	5	29
P_03	34	0	8	14	10	30
P_04	17	0	15	4	11	27
P_06	11	0	6	6	4	30
P_07	20	0	8	0	0	25
P_08	28	0	12	8	12	26
P_10	8	3	2	3	4	28
P_11	14	12	17	19	15	24
P_12	9	0	10	17	18	28
P_13	19	28	5	4	9	27
P_14	31	0	7	11	15	26
P_15	13	12	8	7	9	25
P_16	21	0	1	4	12	25
P_17	12	0	0	2	4	27
P_18	17	11	5	12	20	24

Table 7-1: clinical data from patients in experiment 9.

Performance

All participants learned the action-outcome contingencies. Reversal events were predictably followed by more errors, but both patients and healthy controls returned to better-than chance performance within a few trials (see figure 7-1).

For each participant, a performance score was calculated, defined as the ratio of error trials to non-error trials across the experiment. An independent samples t-test of these ratios showed that patients chose the non-rewarded key more often than controls (patients: M = 1.00, SD = .11, controls: M = .83, SD = .13; $t_{(31)} = 4.17$, $p < .001$)

Action Binding

Action binding results were submitted to a 2x2 mixed ANOVA with a between-subjects factor of group (PD or controls) and a within-subjects factor of outcome on the previous trial (rewarded or penalised). There was no main effect of group ($F_{(1,31)} = 2.01$, $p = .166$, $\eta_p^2 = .061$) and no main effect of previous outcome ($F_{(1,31)} = .064$, $p = .802$, $\eta_p^2 = .002$). However there was a significant interaction ($F_{(1,31)} = 5.693$, $p = 0.023$, $\eta_p^2 = .155$) (see figure 7-1). To explore the interaction, binding following a positive outcome was subtracted from binding following a negative outcome. The result is a measure of the post-error action boost. Healthy participants showed significantly more of this boost than patients ($t_{(31)} = 2.386$, $p = .023$). Indeed, the normal pattern of increased binding following error was reversed in the patient group.

An independent-samples t-test of switching frequency revealed no effect of group ($t_{(31)} = -1.637$, $p = .112$). There is therefore no evidence that patients were more likely to perseverate actions than healthy controls.

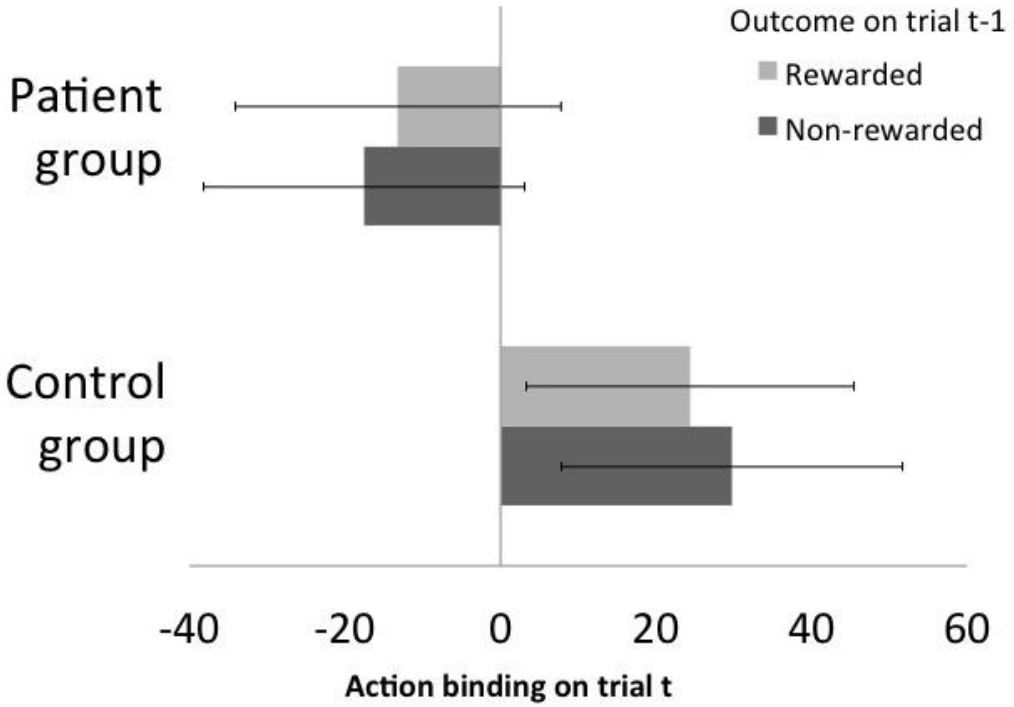

Figure 7-1: Action binding in experiment 9. Error bars represent standard errors.

Modelling & learning rate analysis

A modelling and learning rate analysis was conducted as in chapter 4 (section 4.5.).

It was predicted that patients would not show the same boost to action binding observed in previous studies with healthy participants.

Data and model are shown in figure 7-2.

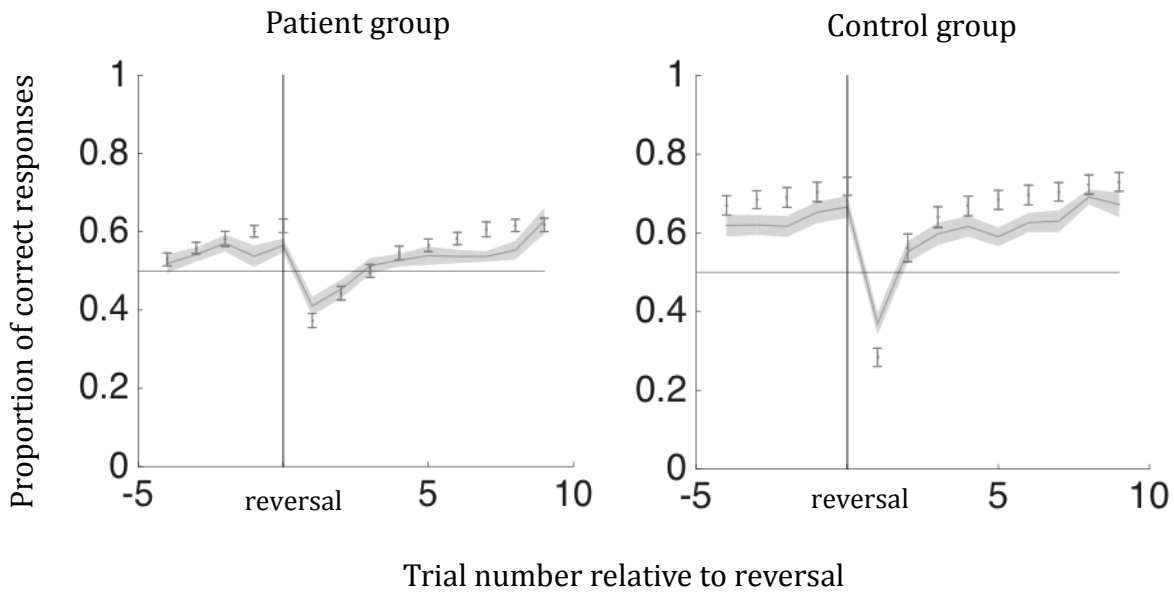

Figure 7-2: Proportion of correct responses before and after a reversal event for patients and healthy controls. The continuous line represents patient data and individual vertical markers represent predictions of the reinforcement-learning model.

To determine if patients and controls displayed differences in learning, the asymmetric learning rates α^+ and α^- were submitted to a mixed ANOVA with a between-subjects factor of group (patient/control) and a within-subjects factor of prediction error sign (positive/ negative).

There was a main effect of group ($F_{(1,32)}$ = 8.41, p = .007, η_p^2 = .213), with controls using prediction errors to update their preferences more than patients. We also found a main effect of prediction error sign ($F_{(1,32)}$ = 21.75, p = .000, η_p^2 = .412). All participants were more influenced by positive than negative prediction errors when updating action preferences. Finally, there was a significant interaction ($F_{(1,32)}$ = 5.6, p = .024, η_p^2 = .153) (see figure 7-3). Simple effects t-tests showed that patients were significantly less likely than controls to update preferences following negative prediction errors ($t_{(31)}$ = 3.44, p = .002) but not positive outcomes ($t_{(31)}$ = .28, p = .97) .

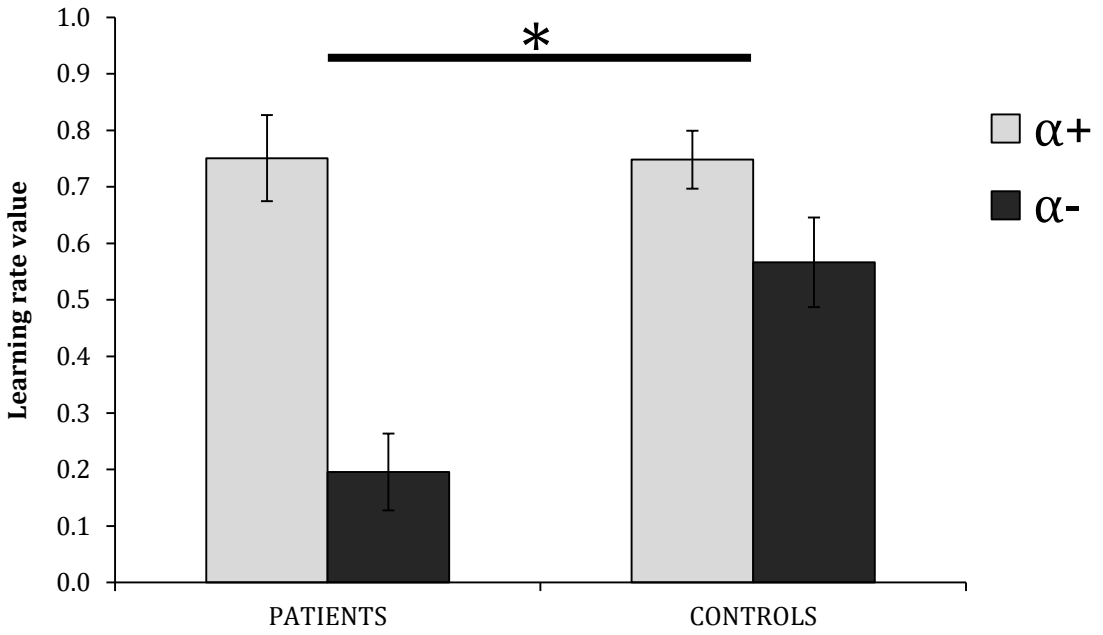

Figure 7-3: Learning rates for patients and controls. The values α^+ and α^- represent, respectively, the degree to which action preferences are updated following a positive or negative prediction error.

7.2.4. Discussion

Parkinson's disease patients on medication and age-matched healthy controls both completed a reversal-learning task with an embedded intentional binding measure. Both groups demonstrated an ability to obtain reward through action selection, although the patient group performed worse overall. A simple reinforcement-learning model showed that patients and controls updated action preferences in response to positive prediction errors. However, the learning rate from negative prediction errors was reduced in patients relative to controls.

Chronometric measures of the perceived time of action obtained during the task were examined for intentional binding effects, and specifically the tendency for the perceived time of an action to shift towards the corresponding outcome. For healthy controls only, trials following an error showed greater action binding than trials following a success,

suggesting a post-error agency boost. Patients with PD showed neither an overall action binding effect, nor a post-error agency boost.

A previous study (Moore et al., 2010) reported stronger intentional binding in a small sample of PD patients ON medication than in healthy controls. This effect was not replicated here. However, the result of Moore and colleagues could be interpreted as a drug effect, rather than a disease effect, since PD patients OFF medication did not differ from the control group. Further, their effect depended on outcome binding as well as action binding, and outcome binding was not measured in this study.

Crucially, the post-error agency boost was reduced in the PD group, as shown by a significant ANOVA interaction between group and previous outcome. Thus, the patients' performance in the probabilistic learning task was characterised by alterations of behavioural adaptation, and also of subjective experience. First, the PD patients used negative prediction errors significantly less than controls when updating action choice preferences. Second, they showed reduced action binding following such errors, relative to healthy volunteers. The results suggest that these deficits may be linked.

A diminished capacity to learn from environmental cues may be a reason for the absence of a PEAB in the patient group. In this sense, patients may be likened to participants in the *difficult* condition of experiment 5, who also did not show increased binding after an error. The observed lack of a PEAB in this experiment is therefore particularly interesting, as it points to a hitherto unknown effect of cognitive impairments on the subjective experience of voluntary actions. Impulsivity in PD may short-circuit meaningful engagement in risky choice tasks, such that errors are not appreciated as valuable information and do not trigger an agency boost. Interestingly, research shows that SoA is inherently rewarding (Karsh & Eitam, 2015). Reduced SoA

might then partially account for known motivational deficits in PD, such as apathy and depression. Future experimental work with adequate sample sizes could determine if these aspects of the pathology are indeed linked.

These results add to previous experimental results indicating cognitive deficits in Parkinson's disease. Cools, Barker, Sahakian, & Robbins (2001) found that patients performed worse in a task-switching paradigm when medication was withdrawn. Interestingly, they found the opposite effect when patients completed a reversal-learning task similar to the paradigm reported here. Performance was worse when patients were medicated (as observed here). These combined results support the theory that an increase in dopamine has a differential effect on cognitive functions that may depend on the activation of separate areas of the prefrontal cortex (Gotham et al., 1988; Swainson, 2000). Since patients in experiment 9 were tested only ON medication, the separate contributions of dopamine-depleting disease, and dopaminergic medication cannot be adequately determined.

The diminished ability to update preferences following errors observed here was associated with, indeed may be the cause of, a reduction SoA. These data therefore extend the findings of Cools et al. (2001), Swainson (2000) and other researchers by providing evidence for a subjective experience that accompanies deficits in cognitive processes in Parkinson's disease.

As an alternative explanation for the absent post-error boost in patients, the influence of action perseveration was also considered. If patients perseverated their actions more than healthy controls, this could potentially have produced an attentional bias towards actions rather than the outcomes they produced. However, there was no statistically significant increase in perseveration in patients.

Another possible explanation for the post-error SoA boost is an attempt to reassert control following an error. Cognitive processes that follow detrimental outcomes include post-error slowing (PES) (Danielmeier & Ullsperger, 2011). Botvinick et al. (2001) suggested that PES might occur due to strategic adjustments and increased caution in responding. As such, one would expect deficits in strategic adjustment in groups such as Parkinson's, where cognitive control is impaired.

The learning bias observed here could be driven by pathology, or by the influence of dopamine through medication. The patient group in this experiment was remarkably homogenous in terms of disease severity and levels of dopaminergic medication. This would need to be investigated in a potential future study testing patients both ON and OFF medication.

7.3. General conclusions from experiment 9

Patients experience a reduced SoA in response to negative outcomes in goal-based learning compared to healthy controls. This effect is not a consequence of action perseveration, but may be due to the diminished influence of negative prediction errors on learning in Parkinson's disease.

The data obtained in chapter 7 complement and expand on the results presented in chapters 2 and 3. The combined results describe the subjective experience of voluntary actions in PD and how this experience is used to guide future actions.

Results from all chapters are considered in the following general discussion.

Chapter 8: General Discussion

8.1. Summary of experimental work in this thesis

This thesis explored human voluntary actions in health and disease. Distinct from reflexes and other automatic movements, voluntary actions are accompanied by a unique subjective experience. This experience, volition, provides what Shadlen calls the 'freedom from immediacy' (Shadlen & Gold, 2004). Further, voluntary actions are a critical component of the sense of agency (SoA), the feeling that we control our actions, and thereby influence changes in the outside world.

The work presented here attempted to further research in the fields of motor control and action awareness with two specific aims:

Firstly, the subjective experience of volition in Parkinson's disease was investigated by comparing cognitive and neurophysiological measures of volition within the *same* patients while OFF and ON two symptom-alleviating interventions.

Secondly, a new paradigm investigating the SoA was introduced, tested and evaluated. This paradigm combined an implicit measure of agency with a goal-based decision-making task, in order to provide an ecologically valid context in which participants could make actions for a reason.

In nine experiments, mental chronometry was used as means of determining awareness of intentions and actions. The Libet paradigm (Libet et al., 1983), intentional binding (Haggard, Clark and Kalogeras, 2002) and reversal learning tasks (Rolls, 1999; Cools, Clark, Owen & Robbins, 2002) were used to investigate the specific factors that influence the subjective experience of volition and agency. Electroencephalography

(EEG) was used to examine the neural processes thought to underlie volition, both in healthy participants and in PD patients. A meta-analysis was used to evaluate the strength of the effects obtained using the novel paradigm introduced in chapter 4. The findings of these experiments are discussed below.

8.2. Volition in Parkinson's disease

Experiment 1 showed that PD patients were able to adequately complete a Libet task without any diminished precision relative to healthy controls. Interestingly, these results differed from Tabu et al. (2015), who found delayed awareness of intentions in PD.

Increasing DA availability through dopaminergic medication resulted in earlier judgements of both action and action intention awareness in PD patients. Alleviation of DA-related symptoms by DBS did not appear to have the same effect. The 'dopamine overdose hypothesis' (Vaillancourt et al., 2013), suggests that earlier judgements may have arisen from patients being in a *hyper*-dopaminergic state.

Interestingly, while medication and DBS implicate different neural mechanisms, a follow-up analysis showed that their common functional impact might have driven earlier judgements in the Libet task. A direct comparison of data from experiments 1 and 2 showed that there was a significant effect of being ON intervention, which did not interact with the specific type of intervention.

Interpreting the results of experiments 2 and 3 using a threshold model would suggest that rate change in the accumulation of motor evidence was not a convincing explanation for earlier awareness of actions and intentions observed while patients

were ON intervention. Rate change would have manifested as a steeper or shallower accumulation curve, resulting in a larger or smaller gap between intention and action awareness. However the results of chapter 2 showed that these timing judgements moved together, so the 'intention-action gap' remained constant.

Experiment 3 provided neurological evidence that the observed changes in intention and action awareness may be due to an earlier onset to the readiness potential, and/or a lowering of perceptual thresholds. Furthermore, an analysis of the variability of amplitude within individual trials suggested that prior to action onset, there was a more marked drop in variability in patients ON intervention than patients OFF intervention. One potential explanation is that while ON intervention, the motor preparation environment may become more efficient (earlier onset) and less noisy (lowered thresholds). These remarks are made with caution however, as the behavioural effect in question was not detected in experiment 3.

Taken together, chapters 2 and 3 provide insight into the neural context in which PD patients generate voluntary actions. The features of this neural context may be associated with earlier perception of intentions and actions when patients are under the influence of symptom-alleviating interventions.

8.3. Sense of Agency: a novel paradigm

Chapter 4 introduced a novel paradigm in which participants were required to learn relations between actions and outcomes. Intentional binding measures were embedded in the task to show how outcomes influence the sense of agency in a dynamic learning environment. This paradigm is seemingly the first to combine techniques from decision-making and motor control research.

Results from several experiments showed that participants were able to complete the task with an appreciation of how their decisions could benefit them. The paradigm therefore represents a useful tool with which to derive implicit measures of agency in a controlled environment where actions have meaning.

The novel paradigm also allowed features of the choice environment to be manipulated, in order to examine their individual effects. Over three experiments, difficulty was manipulated by adjusting the predictability of outcomes. In experiment 6, difficulty was rendered irrelevant altogether in one condition, in which outcomes occurred at random and learning was impossible. The ability to learn proved to be a critical factor in the emergence of a boost to action binding following negative feedback (discussed in the following section).

In experiment 5, the post-error boost was not detected in versions of the task that were made considerably easier, or considerably more difficult. This pattern of results suggests that a learning context alone is insufficient to give rise to the effect. Rather, it may be critical that participants feel that they are able to adequately detect shifting patterns of probability, without these patterns becoming too predictable. The parameters of experiments 4 and 6 might represent a degree of difficulty that induces engagement in the learning aspect of the task without being either too discouraging or too repetitive. This cognitive engagement may be the critical ingredient that gives rise to the PEAB.

Contrary to previous work using valenced outcomes in a binding task (Takahata et al., 2012; Yoshie & Haggard, 2013), experiment 4 found an *increase* in outcome binding on trials with negative outcomes. Outcome binding was not explored further in this thesis, however the most obvious difference between experiment 4 and the cited work is the

ability to learn from one trial to the next. Future experimental research may therefore test whether the ability to learn may influence the effect of valenced outcomes on different components of intentional binding.

8.4. Sense of Agency: Post-error Action boost

The experimental results of chapters 4 and 5 showed that actions following negative feedback were more temporally bound towards their outcomes than actions following positive feedback. This effect suggests that when people experience negative outcomes, they feel more control, not less, when they make subsequent related actions.

Research previously suggested that binding is insensitive to action-outcome relationships (Desantis, Hughes, & Waszak, 2012). Evidence for the reliability of the post-error action boost (PEAB) was therefore obtained using statistical meta-analysis in chapter 6. The effect proved robust, and emerged as a small, but consistent effect across the included studies.

The PEAB may seem counterintuitive, however the experimental design encouraged participants to make actions in pursuit of goals, and evaluate whether their actions were successful or not. An analysis of learning rates in experiments 4 and 6 confirmed that participants did engage with the task in this way. Negative feedback might have increased participants' feeling of agency for a short time to facilitate learning. It was through failures that participants learned to act appropriately in the future.

Intentional binding is a proxy measure of agency, so results must always be interpreted with caution. Binding may be influenced by factors other than SoA. Most alternative theories propose that binding merely represents an appreciation of causation generally

(Buehner & Humphreys, 2009; Buehner (2012), Cravo, Craessens, & Baldo, 2009; Cravo, Craessens, and Baldo, 2011). However, the experimental designs presented in this work always compared data between critical conditions, in order to isolate individual influences on binding measures. The emergence of the PEAB is difficult to reconcile with theories of general causation, as causation is presumably equivalent regardless of the type of outcome. However, inferring agency from binding always remains a critical assumption and must be acknowledged as such.

Interpreting the PEAB in the context of learning and decision-making is also inherently limited by the methods used. The timescales used in Libet tasks are short, typically less than a second, and therefore preclude strong assertions about decision-making over longer timescales. Choosing which button to press in an experimental task cannot be easily compared to selecting an appropriate outfit or buying a home.

Nevertheless, the findings presented here are noteworthy for two reasons. Firstly, the PEAB was proven to be reliable and replicable. Future experimental work on sense of agency, learning, causation and decision-making can therefore make use of this effect in experimental design. The PEAB is ideally suited to provide an index of engagement in learning tasks, alongside other currently used metrics such as performance data and eye-tracking.

Secondly, these results offer, seemingly for the first time, an explanation of *why* agency manifests. Agency research has hitherto been preoccupied with establishing reliable measures and defining putative mechanisms (Haggard, 2017). However, the present work suggests that SoA might support the motivation to improve performance. The brain appears to house a cognitive/experiential mechanism that ensures occasional failures do not immediately prevent further exploration. Indeed, SoA appears to

increase precisely when cognitive engagement is required to distinguish incidental failures from systematically detrimental outcomes. Experiments in this thesis therefore break new ground in linking the subjective experience of agency to the cognitive mechanisms of reinforcement learning, and thereby offering a reason why agency should exist at all.

8.5. Sense of Agency in Parkinson's disease

Chapter 7 investigated the PEAB in a PD population ON medication. The known cognitive deficits of PD (deVos et al., 1996; Weintraub, David, Evans, Grant, & Stacy, 2015) were expected to influence the learning-related effect seen in healthy participants.

In experiment 9, PD patients, unlike healthy controls, did not show a PEAB. Analysis of choice behaviour showed that this effect was not simply a consequence of action perseveration. Rather, a learning rate analysis demonstrated that a reduced PEAB in patients might be attributable to a lesser influence of negative prediction errors on learning, compared to controls.

This experiment further validated the use of the novel paradigm, while expanding on the data obtained in chapters 2 & 3. The combined results provide greater perspective on the perception of voluntary actions in Parkinson's disease, including how the outcomes of such actions might be evaluated.

8.6. Further directions

A comprehensive understanding of voluntary actions in PD remains elusive. More research is required to determine the influence and interactions of therapeutic interventions, comorbid cognitive deficits and other factors, such as age, on the subjective experience of making and monitoring voluntary actions in PD. The effects described in this thesis require replication in larger subject groups, to more adequately assess effect size and robustness.

The post-error action boost is inextricably linked to voluntary actions, yet the precise neural underpinnings of agency are still unclear, particularly in a more complex decision-making environment. The novel paradigm presented here has proven itself as a reliable cognitive task that can be readily modified. It might therefore be used to design experiments that elicit a post-error action boost with a view to measuring the spatial and temporal nature of neural activity in the premotor cortex and beyond.

8.7. Conclusions

This thesis examined the subjective experience of initiating and monitoring voluntary actions in health and disease. Nine experiments and a meta-analysis investigated changes in time perception that reflect underlying cognitive and neurological processes. The use of patient groups, electrophysiology and a novel behavioural paradigm led to new insights being discovered, replicated and evaluated.

This work has described mechanisms that contribute to learning in humans, both in health and disease. When faced with unexpected outcomes, the brain ensures that learning persists and informs our goal-directed movements. When physical movement is impaired, the use of symptom-alleviating interventions appears to influence not only motor function, but the subjective experience that accompanies purposeful actions.

This thesis represents a convergence of several research topics for the first time. The experience of voluntary actions was found to be associated with dopaminergic drives in the premotor cortex, and the sense of agency was investigated in the context of learning and decision-making. These convergences represent fruitful avenues of investigation for future research on motor control, voluntary actions, decision-making and movement-related pathology and treatment.

The overall picture is one of adaptability: as humans we use learning to refine and adjust our behaviour. The work presented here highlights the role of voluntary actions in learning to adapt. Ultimately, it is this ability to respond intelligently to rapidly changing environments that has ensured the on-going evolution of human society and the continued development of the manmade world.

9. Appendices

9.1 Appendix I: clinical instruments used in chapters 2, 3 and 7.

Figure 9-1 : Unified Parkinson's Disease Rating Scale Part 3 motor examination (UPDRS-III)

18. Speech

0 = Normal.
1 = Slight loss of expression, diction and/or volume.
2 = Monotone, slurred but understandable; moderately impaired.
3 = Marked impairment, difficult to understand.
4 = Unintelligible.

19. Facial Expression

0 = Normal.
1 = Minimal hypomimia, could be normal "Poker Face".
2 = Slight but definitely abnormal diminution of facial expression
3 = Moderate hypomimia; lips parted some of the time.
4 = Masked or fixed facies with severe or complete loss of facial expression; lips parted 1/4 inch or more.

20. Tremor at rest (head, upper and lower extremities)

0 = Absent.
1 = Slight and infrequently present.
2 = Mild in amplitude and persistent. Or moderate in amplitude, but only intermittently present.
3 = Moderate in amplitude and present most of the time.
4 = Marked in amplitude and present most of the time.

21. Action or Postural Tremor of hands

0 = Absent.
1 = Slight; present with action.
2 = Moderate in amplitude, present with action.
3 = Moderate in amplitude with posture holding as well as action.
4 = Marked in amplitude; interferes with feeding.

22. Rigidity (Judged on passive movement of major joints with patient relaxed in sitting position. Cogwheeling to be ignored.)

0 = Absent.
1 = Slight or detectable only when activated by mirror or other movements.
2 = Mild to moderate.
3 = Marked, but full range of motion easily achieved.
4 = Severe, range of motion achieved with difficulty.

23. Finger Taps (Patient taps thumb with index finger in rapid succession.)

0 = Normal.
1 = Mild slowing and/or reduction in amplitude.
2 = Moderately impaired. Definite and early fatiguing. May have occasional arrests in movement.
3 = Severely impaired. Frequent hesitation in initiating movements or arrests in ongoing movement.
4 = Can barely perform the task.

24. Hand Movements (Patient opens and closes hands in rapid succesion.)

0 = Normal.
1 = Mild slowing and/or reduction in amplitude.

2 = Moderately impaired. Definite and early fatiguing. May have occasional arrests in movement.
3 = Severely impaired. Frequent hesitation in initiating movements or arrests in ongoing movement.
4 = Can barely perform the task.

25. Rapid Alternating Movements of Hands (Pronation-supination movements of hands, vertically and horizontally, with as large an amplitude as possible, both hands simultaneously.)

0 = Normal.
1 = Mild slowing and/or reduction in amplitude.
2 = Moderately impaired. Definite and early fatiguing. May have occasional arrests in movement.
3 = Severely impaired. Frequent hesitation in initiating movements or arrests in ongoing movement.
4 = Can barely perform the task.

26. Leg Agility (Patient taps heel on the ground in rapid succession picking up entire leg. Amplitude should be at least 3 inches.)

0 = Normal.
1 = Mild slowing and/or reduction in amplitude.
2 = Moderately impaired. Definite and early fatiguing. May have occasional arrests in movement.
3 = Severely impaired. Frequent hesitation in initiating movements or arrests in ongoing movement.
4 = Can barely perform the task.

27. Arising from Chair (Patient attempts to rise from a straightbacked chair, with arms folded across chest.)

0 = Normal.
1 = Slow; or may need more than one attempt.
2 = Pushes self up from arms of seat.
3 = Tends to fall back and may have to try more than one time, but can get up without help.
4 = Unable to arise without help.

28. Posture

0 = Normal erect.
1 = Not quite erect, slightly stooped posture; could be normal for older person.
2 = Moderately stooped posture, definitely abnormal; can be slightly leaning to one side.
3 = Severely stooped posture with kyphosis; can be moderately leaning to one side.
4 = Marked flexion with extreme abnormality of posture.

29. Gait

0 = Normal.
1 = Walks slowly, may shuffle with short steps, but no festination (hastening steps) or propulsion.
2 = Walks with difficulty, but requires little or no assistance; may have some festination, short steps, or propulsion.
3 = Severe disturbance of gait, requiring assistance.
4 = Cannot walk at all, even with assistance.

30. Postural Stability (Response to sudden, strong posterior displacement - pull on shoulders while patient erect with eyes open and feet slightly apart. Patient is prepared.)

0 = Normal.
1 = Retropulsion, but recovers unaided.
2 = Absence of postural response; would fall if not caught by examiner.

3 = Very unstable, tends to lose balance spontaneously.
4 = Unable to stand without assistance.

31. Body Bradykinesia and Hypokinesia (Combining slowness, hesitancy, decreased armswing, small amplitude, and poverty of movement in general.)

0 = None.
1 = Minimal slowness, giving movement a deliberate character; could be normal for some persons. Possibly reduced amplitude.
2 = Mild degree of slowness and poverty of movement which is definitely abnormal. Alternatively, some reduced amplitude.
3 = Moderate slowness, poverty or small amplitude of movement.
4 = Marked slowness, poverty or small amplitude of movement.

Figure 9-2: Montreal Cognitive Assessment

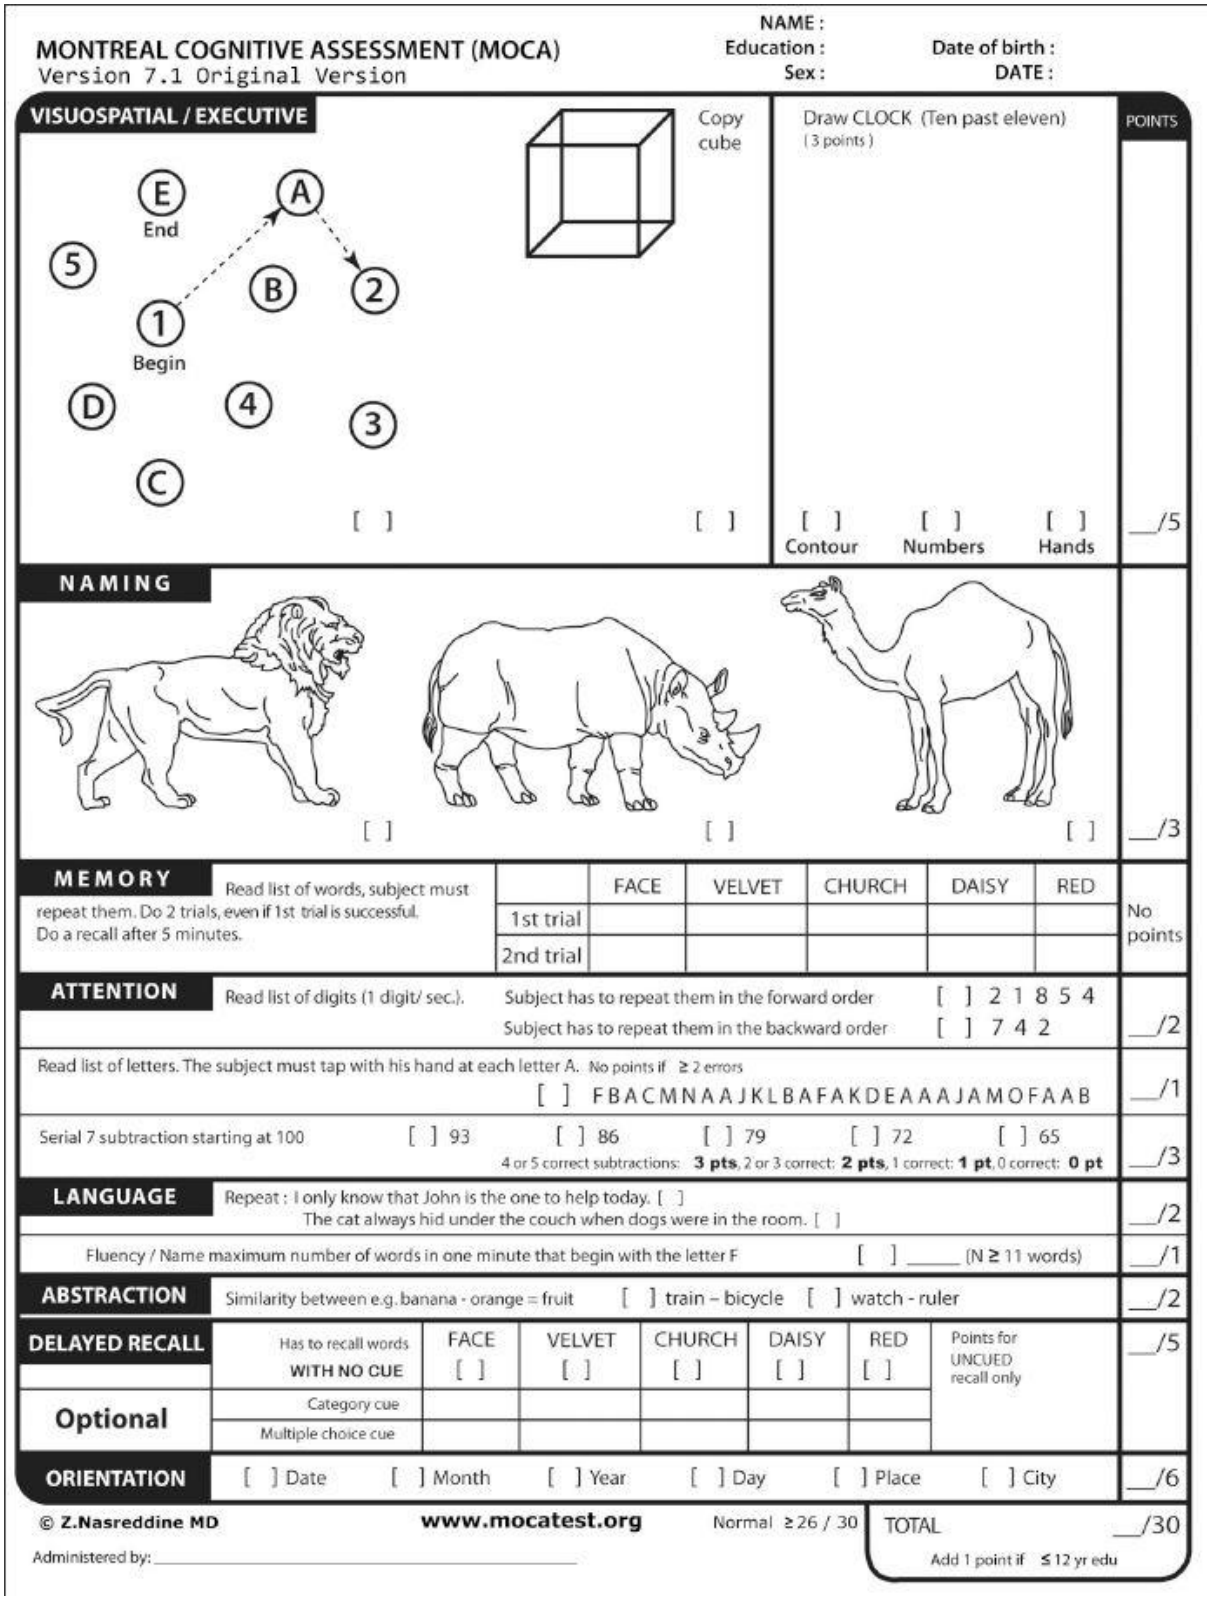

MONTREAL COGNITIVE ASSESSMENT (MOCA)
Version 7.1 Original Version

NAME :
Education : Date of birth :
Sex : DATE :

VISUOSPATIAL / EXECUTIVE — POINTS

Copy cube [ ] [ ]

Draw CLOCK (Ten past eleven) (3 points) [ ] Contour [ ] Numbers [ ] Hands __/5

NAMING __/3

[ ] [ ] [ ]

MEMORY Read list of words, subject must repeat them. Do 2 trials, even if 1st trial is successful. Do a recall after 5 minutes.

	FACE	VELVET	CHURCH	DAISY	RED
1st trial					
2nd trial					

No points

ATTENTION Read list of digits (1 digit/ sec.).
Subject has to repeat them in the forward order [ ] 2 1 8 5 4
Subject has to repeat them in the backward order [ ] 7 4 2 __/2

Read list of letters. The subject must tap with his hand at each letter A. No points if ≥ 2 errors
[ ] F B A C M N A A J K L B A F A K D E A A A J A M O F A A B __/1

Serial 7 subtraction starting at 100 [ ] 93 [ ] 86 [ ] 79 [ ] 72 [ ] 65
4 or 5 correct subtractions: **3 pts**, 2 or 3 correct: **2 pts**, 1 correct: **1 pt**, 0 correct: **0 pt** __/3

LANGUAGE Repeat : I only know that John is the one to help today. [ ]
The cat always hid under the couch when dogs were in the room. [ ] __/2

Fluency / Name maximum number of words in one minute that begin with the letter F [ ] _____ (N ≥ 11 words) __/1

ABSTRACTION Similarity between e.g. banana - orange = fruit [ ] train – bicycle [ ] watch - ruler __/2

DELAYED RECALL __/5

	Has to recall words WITH NO CUE	FACE [ ]	VELVET [ ]	CHURCH [ ]	DAISY [ ]	RED [ ]	Points for UNCUED recall only
Optional	Category cue						
	Multiple choice cue						

ORIENTATION [ ] Date [ ] Month [ ] Year [ ] Day [ ] Place [ ] City __/6

© Z.Nasreddine MD www.mocatest.org Normal ≥ 26 / 30

TOTAL __/30
Add 1 point if ≤ 12 yr edu

Administered by: _____

Figure 9-3: Questionnaire for Impulsive-compulsive Behaviour Disorders in PD (QUIP-RS)

Questionnaire for Impulsive-Compulsive Disorders in Parkinson's Disease - Rating Scale (QUIP-RS)

Reported by: _____ Patient _____ Informant _____ Patient and Informant

Patient / Subject: _____

Date: _____

1. How much do you think about the following behaviors (such as having trouble keeping thoughts out of your mind or feeling guilty)?

Gambling?	___Never(0)	___Rarely(1)	___Sometimes(2)	___Often(3)	___Very often(4)
Sex?	Never(0)	Rarely(1)	Sometimes(2)	Often(3)	Very often(4)
Buying?	___Never(0)	___Rarely(1)	___Sometimes(2)	___Often(3)	___Very often(4)
Eating?	___Never(0)	___Rarely(1)	___Sometimes(2)	___Often(3)	___Very often(4)
Performing tasks or hobbies?	___Never(0)	___Rarely(1)	___Sometimes(2)	___Often(3)	___Very often(4)
Repeating simple activities?	___Never(0)	___Rarely(1)	___Sometimes(2)	___Often(3)	___Very often(4)
Taking your PD medications?	___Never(0)	___Rarely(1)	___Sometimes(2)	___Often(3)	___Very often(4)

2. Do you have urges or desires for the following behaviors that you feel are excessive or cause you distress (including becoming restless or irritable when unable to participate in them)?

Gambling?	___Never(0)	___Rarely(1)	___Sometimes(2)	___Often(3)	___Very often(4)
Sex?	Never(0)	Rarely(1)	Sometimes(2)	Often(3)	Very often(4)
Buying?	___Never(0)	___Rarely(1)	___Sometimes(2)	___Often(3)	___Very often(4)
Eating?	___Never(0)	___Rarely(1)	___Sometimes(2)	___Often(3)	___Very often(4)
Performing tasks or hobbies?	___Never(0)	___Rarely(1)	___Sometimes(2)	___Often(3)	___Very often(4)
Repeating simple activities?	___Never(0)	___Rarely(1)	___Sometimes(2)	___Often(3)	___Very often(4)
Taking your PD medications?	___Never(0)	___Rarely(1)	___Sometimes(2)	___Often(3)	___Very often(4)

3. Do you have difficulty controlling the following behaviors (such as increasing them over time, or having trouble cutting down or stopping them)?

Gambling?	___Never(0)	___Rarely(1)	___Sometimes(2)	___Often(3)	___Very often(4)
Sex?	Never(0)	Rarely(1)	Sometimes(2)	Often(3)	Very often(4)
Buying?	___Never(0)	___Rarely(1)	___Sometimes(2)	___Often(3)	___Very often(4)
Eating?	___Never(0)	___Rarely(1)	___Sometimes(2)	___Often(3)	___Very often(4)
Performing tasks or hobbies?	___Never(0)	___Rarely(1)	___Sometimes(2)	___Often(3)	___Very often(4)
Repeating simple activities?	___Never(0)	___Rarely(1)	___Sometimes(2)	___Often(3)	___Very often(4)
Taking your PD medications?	___Never(0)	___Rarely(1)	___Sometimes(2)	___Often(3)	___Very often(4)

4. Do you engage in activities specifically to continue the following behaviors (such as hiding what you are doing, lying, hoarding things, borrowing from others, accumulating debt, stealing, or being involved in illegal acts)?

Gambling?	___Never(0)	___Rarely(1)	___Sometimes(2)	___Often(3)	___Very often(4)
Sex?	Never(0)	Rarely(1)	Sometimes(2)	Often(3)	Very often(4)
Buying?	___Never(0)	___Rarely(1)	___Sometimes(2)	___Often(3)	___Very often(4)
Eating?	___Never(0)	___Rarely(1)	___Sometimes(2)	___Often(3)	___Very often(4)
Performing tasks or hobbies?	___Never(0)	___Rarely(1)	___Sometimes(2)	___Often(3)	___Very often(4)
Repeating simple activities?	___Never(0)	___Rarely(1)	___Sometimes(2)	___Often(3)	___Very often(4)
Taking your PD medications?	___Never(0)	___Rarely(1)	___Sometimes(2)	___Often(3)	___Very often(4)

Figure 9-4: Beck's Depression Inventory

Beck's Depression Inventory

This depression inventory can be self-scored. The scoring scale is at the end of the questionnaire.

1.
- 0 I do not feel sad.
- 1 I feel sad
- 2 I am sad all the time and I can't snap out of it.
- 3 I am so sad and unhappy that I can't stand it.

2.
- 0 I am not particularly discouraged about the future.
- 1 I feel discouraged about the future.
- 2 I feel I have nothing to look forward to.
- 3 I feel the future is hopeless and that things cannot improve.

3.
- 0 I do not feel like a failure.
- 1 I feel I have failed more than the average person.
- 2 As I look back on my life, all I can see is a lot of failures.
- 3 I feel I am a complete failure as a person.

4.
- 0 I get as much satisfaction out of things as I used to.
- 1 I don't enjoy things the way I used to.
- 2 I don't get real satisfaction out of anything anymore.
- 3 I am dissatisfied or bored with everything.

5.
- 0 I don't feel particularly guilty
- 1 I feel guilty a good part of the time.
- 2 I feel quite guilty most of the time.
- 3 I feel guilty all of the time.

6.
- 0 I don't feel I am being punished.
- 1 I feel I may be punished.
- 2 I expect to be punished.
- 3 I feel I am being punished.

7.
- 0 I don't feel disappointed in myself.
- 1 I am disappointed in myself.
- 2 I am disgusted with myself.
- 3 I hate myself.

8.
- 0 I don't feel I am any worse than anybody else.
- 1 I am critical of myself for my weaknesses or mistakes.
- 2 I blame myself all the time for my faults.
- 3 I blame myself for everything bad that happens.

9.
- 0 I don't have any thoughts of killing myself.
- 1 I have thoughts of killing myself, but I would not carry them out.
- 2 I would like to kill myself.
- 3 I would kill myself if I had the chance.

10.
- 0 I don't cry any more than usual.
- 1 I cry more now than I used to.
- 2 I cry all the time now.
- 3 I used to be able to cry, but now I can't cry even though I want to.

11.

0 I am no more irritated by things than I ever was.
1 I am slightly more irritated now than usual.
2 I am quite annoyed or irritated a good deal of the time.
3 I feel irritated all the time.

12.

0 I have not lost interest in other people.
1 I am less interested in other people than I used to be.
2 I have lost most of my interest in other people.
3 I have lost all of my interest in other people.

13.

0 I make decisions about as well as I ever could.
1 I put off making decisions more than I used to.
2 I have greater difficulty in making decisions more than I used to.
3 I can't make decisions at all anymore.

14.

0 I don't feel that I look any worse than I used to.
1 I am worried that I am looking old or unattractive.
2 I feel there are permanent changes in my appearance that make me look unattractive
3 I believe that I look ugly.

15.

0 I can work about as well as before.
1 It takes an extra effort to get started at doing something.
2 I have to push myself very hard to do anything.
3 I can't do any work at all.

16.

0 I can sleep as well as usual.
1 I don't sleep as well as I used to.
2 I wake up 1-2 hours earlier than usual and find it hard to get back to sleep.
3 I wake up several hours earlier than I used to and cannot get back to sleep.

17.

0 I don't get more tired than usual.
1 I get tired more easily than I used to.
2 I get tired from doing almost anything.
3 I am too tired to do anything.

18.

0 My appetite is no worse than usual.
1 My appetite is not as good as it used to be.
2 My appetite is much worse now.
3 I have no appetite at all anymore.

19.

0 I haven't lost much weight, if any, lately.
1 I have lost more than five pounds.
2 I have lost more than ten pounds.
3 I have lost more than fifteen pounds.

20.

0 I am no more worried about my health than usual.
1 I am worried about physical problems like aches, pains, upset stomach, or constipation.
2 I am very worried about physical problems and it's hard to think of much else.
3 I am so worried about my physical problems that I cannot think of anything else.

21.

0 I have not noticed any recent change in my interest in sex.
1 I am less interested in sex than I used to be.
2 I have almost no interest in sex.
3 I have lost interest in sex completely.

INTERPRETING THE BECK DEPRESSION INVENTORY

Now that you have completed the questionnaire, add up the score for each of the twenty-one questions by counting the number to the right of each question you marked. The highest possible total for the whole test would be sixty-three. This would mean you circled number three on all twenty-one questions. Since the lowest possible score for each question is zero, the lowest possible score for the test would be zero. This would mean you circles zero on each question. You can evaluate your depression according to the Table below.

Total Score	Levels of Depression
1-10	These ups and downs are considered normal
11-16	Mild mood disturbance
17-20	Borderline clinical depression
21-30	Moderate depression
31-40	Severe depression
over 40	Extreme depression

9.2 Appendix II

Figure 9-5 : Observed 6 Hz artefact from DBS stimulator device in experiment 3

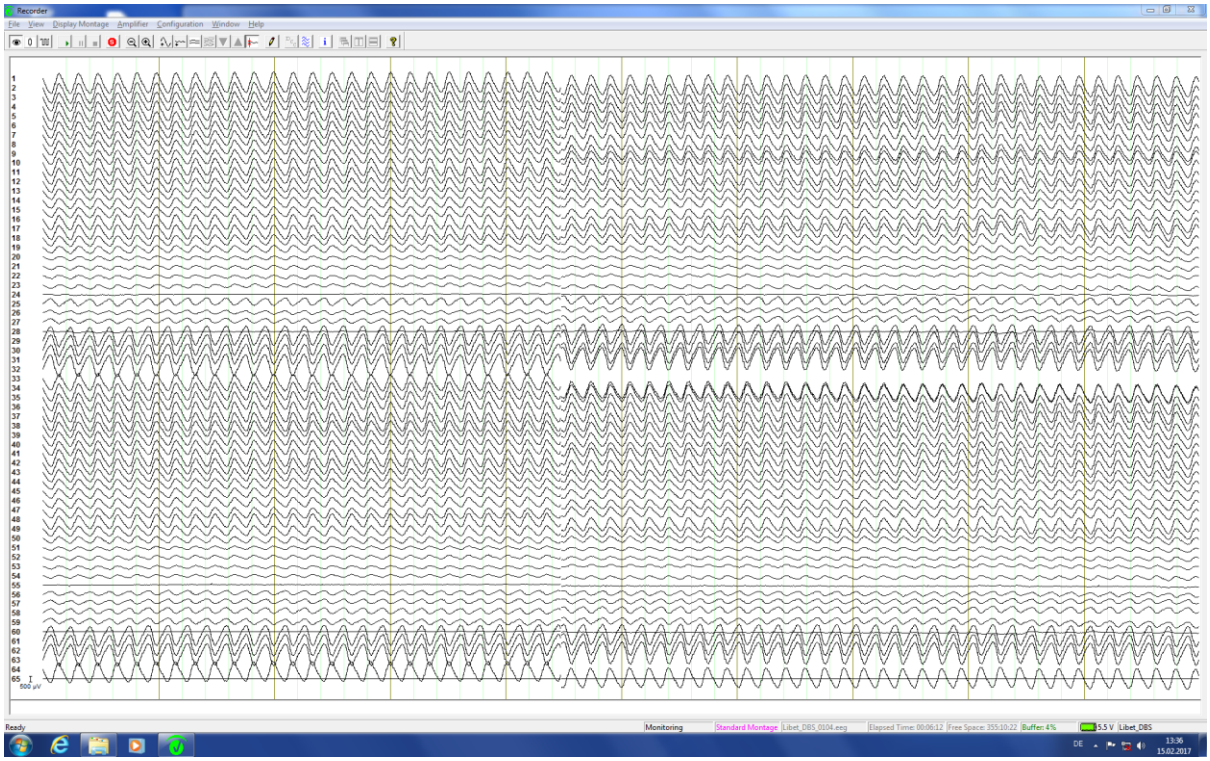

10. References

Akkal, D., Dum, R. P., & Strick, P. L. (2007). Supplementary motor area and presupplementary motor area: targets of basal ganglia and cerebellar output. *Journal of Neuroscience, 27*(40), 10659-10673.

Alloy, L. B., & Abramson, L. Y. (1979). Judgment of contingency in depressed and nondepressed students: Sadder but wiser*?. Journal of experimental psychology: General, 108*(4), 441.

Artieda, J., Pastor, M. A., Lacruz, F., & Obeso, J. A. (1992). Temporal discrimination is abnormal in Parkinson's disease. *Brain, 115*(1), 199-210.

Bandura, A. (1989). Human agency in social cognitive theory. *American psychologist, 44*(9), 1175.

Barrett, G., Shibasaki, H., & Neshige, R. (1986). Cortical potential shifts preceding voluntary movement are normal in parkinsonism. *Electroencephalography and clinical neurophysiology, 63*(4), 340-348.

Banks, W. P., & Pockett, S. (2007). Benjamin Libet's work on the neuroscience of free will. *The Blackwell companion to consciousness*, 657-670.

Bayer, H. M., & Glimcher, P. W. (2005). Midbrain dopamine neurons encode a quantitative reward prediction error signal. *Neuron, 47*(1), 129-141.

Bishop, C. M. (2006). *Pattern recognition and machine learning* (Vol. 4, No. 4, p. 12). New York: springer.

Blakemore, S., Wolpert, D. & Frith, C. (2002). Abnormalities in the awareness of action. *Trends Cognitive Science. 6*, 237–242 (2002).

Bobzien, Susanne (2000). Did Epicurus discover the Free-Will Problem? *Oxford Studies in Ancient Philosophy 19*:287-337.

Botvinick, M. M., Braver, T. S., Barch, D. M., Carter, C. S., & Cohen, J. D. (2001). Conflict monitoring and cognitive control. *Psychological review, 108*(3), 624.

Braak, H., & Braak, E. (2000). Pathoanatomy of Parkinson's disease. *Journal of neurology*, 247, II3-II10.

Braak, H., Braak, E., Yilmazer, D., De Vos, R. A. I., Jansen E. N. H., Bohl J., Jellinger K. (1994) Amygdala pathology in Parkinson's disease. *Acta Neuropathologica 88*:493–500

Braak, H., Braak, E., Yilmazer, D., De Vos, R. A. I., Jansen, E. N. H., & Bohl, J. (1996). Pattern of brain destruction in Parkinson's and Alzheimer's diseases. *Journal of neural transmission, 103*(4), 455-490.

Buehner, M. J. (2012). Understanding the past, predicting the future: causation, not intentional action, is the root of temporal binding. *Psychological science, 23*(12), 1490-1497.

Buehner, M. J., & Humphreys, G. R. (2009). Causal Binding of Actions to Their Effects. *Psychological Science*, *20*(10), 1221–1228.

Ceballos-Baumann, A. O., Boecker, H., Bartenstein, P., von Falkenhayn, I., Riescher, H., Conrad, B., ... & Alesch, F. (1999). A positron emission tomographic study of subthalamic nucleus stimulation in Parkinson disease: enhanced movement-related activity of motor-association cortex and decreased motor cortex resting activity. *Archives of neurology, 56(*8), 997-1003.

Cools, R., Barker, R. A., Sahakian, B. J., & Robbins, T. W. (2001). Enhanced or impaired cognitive function in Parkinson's disease as a function of dopaminergic medication and task demands. *Cerebral cortex, 11*(12), 1136-1143.

Cools, R., Clark, L., Owen, A. M., & Robbins, T. W. (2002). Defining the neural mechanisms of probabilistic reversal learning using event-related functional magnetic resonance imaging. *The Journal of Neuroscience*, *22*(11), 4563–4567.

Coulson, S. E., O'dwyer, N. J., Adams, R. D., & Croxson, G. R. (2004). Expression of emotion and quality of life after facial nerve paralysis. *Otology & neurotology*, 25(6), 1014-1019.

Cravo, A. M., Claessens, P. M. E., & Baldo, M. V. C. (2009). Voluntary action and causality in temporal binding. *Experimental Brain Research, 199*(1), 95–99.

Cravo, A. M., Claessens, P. M. E., & Baldo, M. V. C. (2011). The relation between action, predictability and temporal contiguity in temporal binding. *Acta Psychologica, 136*(1), 157–166.

Cui R. & Deecke L. (1999) High resokution DC-EEG analysis of the Bereitschaftspotential and post movement onset potentials accompanying uniorbilateral voluntary finger movements. *Brain Topography, 11,* 233–49.

Cummings J. L. (1992) Depression and Parkinson's disease: a review. *American Journal of Psychiatry. 149*(4):443–454.

Danielmeier, C., & Ullsperger, M. (2011). Post-error adjustments. *Frontiers in psychology, 2,* 233.

Daw, N. D. (2011). Trial-by-trial data analysis using computational models. *Decision making, affect, and learning: Attention and performance XXIII*, *23*, 3-38.

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

Delorme, A., & Makeig, S. (2004). EEGLAB: an open source toolbox for analysis of singletrial EEG dynamics including independent component analysis. *Journal of Neuroscience Methods, 134*(1), 9–21.

De Havas, J., Ghosh, A., Gomi, H., & Haggard, P. (2016). Voluntary motor commands reveal awareness and control of involuntary movement. *Cognition, 155*, 155-167.

Desantis, A., Hughes, G., & Waszak, F. (2012). Intentional binding is driven by the mere presence of an action and not by motor prediction. *PLoS One, 7*(1), e29557.

Dewey, J. A., & Knoblich, G. (2014). Do implicit and explicit measures of the sense of agency measure the same thing?. *PloS one, 9*(10), e110118.

Dick, J. P. R., Cantello, R., Buruma, O., Gioux, M., Benecke, R., Day, B. L., ... & Marsden, C. D. (1987). The Bereitschaftspotential, L-DOPA and Parkinson's disease. *Electroencephalography and clinical neurophysiology, 66*(3), 263-274.

Drevets, W. C. (2001). Neuroimaging and neuropathological studies of depression: implications for the cognitive-emotional features of mood disorders. *Current opinion in neurobiology, 11*(2), 240-249.

Dutilh, G., Vandekerckhove, J., Forstmann, B. U., Keuleers, E., Brysbaert, M., & Wagenmakers, E. J. (2012). Testing theories of post-error slowing. *Attention, Perception, & Psychophysics, 74*(2), 454-465.

Engbert, K., & Wohlschläger, A. (2007). Intentions and expectations in temporal binding. *Consciousness and Cognition, 16*(2), 255-264.

Engbert, K., Wohlschläger, A., Thomas, R., & Haggard, P. (2007). Agency, subjective time, and other minds. *Journal of Experimental Psychology: Human Perception and Performance*, *33*(6), 1261.

Ersche, K. D., Roiser, J. P., Robbins, T. W., & Sahakian, B. J. (2008). Chronic cocaine but not chronic amphetamine use is associated with perseverative responding in humans. *Psychopharmacology, 197*(3), 421-431.

Ernst, M. O., & Banks, M. S. (2002). Humans integrate visual and haptic information in a statistically optimal fashion. *Nature, 415*(6870), 429-433.

Feve, A. P., Bathien, N., & Rondot, P. (1992). Chronic administration of L-dopa affects the movement-related cortical potentials of patients with Parkinson's disease. *Clinical neuropharmacology, 15*(2), 100-108.

Finger, S. (2001). *Origins of neuroscience: a history of explorations into brain function*. Oxford University Press, USA.

Francis, B. A., & Wonham, W. M. (1976). The internal model principle of control theory. *Automatica, 12*(5), 457-465.

Frank, M. J., Samanta, J., Moustafa, A. A., & Sherman, S. J. (2007). Hold your horses: impulsivity, deep brain stimulation, and medication in parkinsonism. *Science, 318*(5854), 1309-1312.

Frank, M. J., Seeberger, L. C., & O'reilly, R. C. (2004). By carrot or by stick: cognitive reinforcement learning in parkinsonism. *Science, 306*(5703), 1940-1943.

Frank, M. J., Woroch, B. S., & Curran, T. (2005). Error-related negativity predicts reinforcement learning and conflict biases. *Neuron*, *47*(4), 495-501.

Fried, I., Mukamel, R., & Kreiman, G. (2011). Internally generated preactivation of single neurons in human medial frontal cortex predicts volition. *Neuron, 69*(3), 548-562.

Frith, C. D., Blakemore, S. J., & Wolpert, D. M. (2000). Explaining the symptoms of schizophrenia: abnormalities in the awareness of action. *Brain Research Reviews, 31*(2), 357-363.

Gallagher, S. (2000). Philosophical conceptions of the self: implications for cognitive science. *Trends in cognitive sciences, 4*(1), 14-21.

Galton, F. (1889). Instrument for determining reaction-time. *Report of the British Association for the Advancement of Science, 5,* 784-785.

Ganos, C., Asmuss, L., Bongert, J., Brandt, V., Münchau, A., & Haggard, P. (2015). Volitional action as perceptual detection: predictors of conscious intention in adolescents with tic disorders. *Cortex, 64*, 47-54.

Garrison, A. H. (1998). The history of the M'Naghten insanity defense and the use of posttraumatic stress disorder as a basis of insanity. *American Journal of Forensic Psychology, 16*, 39-88.

Georgiev, D., Lange, F., Seer, C., Kopp, B., & Jahanshahi, M. (2016). Movement-related potentials in Parkinson's disease. *Clinical Neurophysiology, 127*(6), 2509-2519.

Goetz, C. G., Tilley, B. C., Shaftman, S. R., Stebbins, G. T., Fahn, S., Martinez‐Martin, P., ... & Dubois, B. (2008). Movement Disorder Society‐sponsored revision of the Unified Parkinson's Disease Rating Scale (MDS‐UPDRS): Scale presentation and clinimetric testing results. *Movement disorders, 23*(15), 2129-2170.

Gotham, A. M., Brown, R. G., & Marsden, C. D. (1988). 'Frontal'cognitive function in patients with Parkinson's disease 'on'and 'off'levodopa. *Brain, 111*(2), 299-321.

Gur, R. E., Kohler, C. G., Ragland, J. D., Siegel, S. J., Lesko, K., Bilker, W. B., & Gur, R. C. (2006). Flat affect in schizophrenia: relation to emotion processing and neurocognitive measures. *Schizophrenia bulletin, 32*(2), 279-287.

Haber, S. N., & Knutson, B. (2010). The reward circuit: linking primate anatomy and human imaging. *Neuropsychopharmacology, 35*(1), 4.

Haggard, P. (2005). Conscious intention and motor cognition. *Trends in cognitive sciences, 9*(6), 290-295.

Haggard, P. (2008). Human volition: towards a neuroscience of will. *Nature Reviews Neuroscience, 9*(12), 934-946.

Haggard, P. (2017). Sense of agency in the human brain. *Nature Reviews Neuroscience, 18*(4), 196-207.

Haggard, P., Clark, S., & Kalogeras, J. (2002). Voluntary action and conscious awareness. *Nature neuroscience*, *5*(4), 382-385.

Haggard, P., & Chambon, V. (2012). Sense of agency. *Current Biology*, *22*(10), R390–R392.

Haggard, P. & Eimer, M. (1999). On the relation between brain potentials and the awareness of voluntary movements. *Experimental Brain Research. 126*, 128–133

Haggard, P., Martin, F., Taylor-Clarke, M., Jeannerod, M., & Franck, N. (2003). Awareness of action in schizophrenia. *Neuroreport, 14*(7), 1081-1085.

Hallett, M. (2007). Volitional control of movement: the physiology of free will. *Clinical Neurophysiology, 118*(6), 1179-1192.

Hedges, L. & Vevea, J. (1998) Fixed- and random-effects models in meta-analysis. *Psychological Methods 3*(486–504).

Higgins, J. P., Thompson, S. G., Deeks, J. J., & Altman, D. G. (2003). Measuring inconsistency in meta-analyses. *BMJ: British Medical Journal, 327*(7414), 557.

Hoehn M, Yahr M (1967). Parkinsonism: onset, progression and mortality. *Neurology*. *17*(5): 427–42.

Hollerman, J. R., & Schultz, W. (1998). Dopamine neurons report an error in the temporal prediction of reward during learning. *Nature neuroscience, 1*(4).

Horvath, J. C., Forte, J. D., & Carter, O. (2015). Quantitative review finds no evidence of cognitive effects in healthy populations from single-session transcranial direct current stimulation (tDCS). *Brain stimulation, 8*(3), 535-550.

Humphreys, G. R., & Buehner, M. J. (2009). Magnitude estimation reveals temporal binding at super-second intervals. Journal of Experimental Psychology: *Human Perception and Performance, 35*(5), 1542.

Ikeda, A., & Shibasaki, H. (2003). Generator mechanisms of Bereitschaftspotentials as studied by epicortical recording in patients with intractable partial epilepsy. In *The Bereitschaftspotential* (pp. 45-59). Springer US.

Jackson, S. R., Parkinson, A., Jung, J., Ryan, S. E., Morgan, P. S., Hollis, C., & Jackson, G. M. (2011). Compensatory neural reorganization in Tourette syndrome*. Current Biology, 21*(7), 580-585.

Jahanshahi, M., Jenkins, I.H., Brown, R.G., Marsden, C.D., Passingham, R.E. & Brooks, D.J. (1995) Self-initiated versus externally triggered movements*, Brain, 118* (Pt4), 913-933.

James, W. (1950). *The principles of psychology, Vol 2*. New York: Dover Publications.

Jellinger, K. A. (1991). Pathology of Parkinson's disease. *Molecular and Chemical Neuropathology, 14*(3), 153-197.

Karsh, N., & Eitam, B. (2015). I control therefore I do: Judgments of agency influence action selection. *Cognition, 138*, 122-131.

Keramati, M., Dezfouli, A., & Piray, P. (2011). Speed/accuracy trade-off between the habitual and the goal-directed processes. *PLoS ComputBiol*, *7*(5), e1002055.

Khalighinejad, N., Di Costa, S., & Haggard, P. (2016). Endogenous action selection processes in dorsolateral prefrontal cortex contribute to sense of agency: a meta-analysis of tDCS studies of 'intentional binding'. *Brain stimulation, 9*(3), 372-379.

Khalighinejad, N., Schurger, A., Desantis, A., Zmigrod, L., & Haggard, P. (2017). Precursor processes of human self-initiated action. *bioRxiv*. https://doi.org/10.1101/120105

Kohnstamm, O. (1915). Demonstration einer katatonieartigen erscheinung beim gesunden (katatonusversuch). *Neurol Zentralbl*, 34, 290-291.

Kornhuber, H. & Deecke, L. (1965). Hirnpotentialänderungen bei Willkürbewegungen und passiven Bewegungen des Menschen: Bereitschaftspotential und reafferente Potentiale. *Pflügers Archiv für die Gesamte Physiologie des Menschen und der Tiere. 284*: 1–17.

Kitamura, J. I., Shibasaki, H., & Kondo, T. (1993). A cortical slow potential is larger before an isolated movement of a single finger than simultaneous movement of two fingers. *Electroencephalography and clinical Neurophysiology, 86*(4), 252-258.

Knowlton, B. J., Mangels, J. A., & Squire, L. R. (1996). A neostriatal habit learning system in humans. *Science, 273*(5280), 1399.

Koller, W. C., Pahwa, R., Lyons, K. E., & Albanese, A. (1999). Surgical treatment of Parkinson's disease. *Journal of the neurological sciences, 167*(1), 1-10.

Kurlan, R. (1997). Diagnostic criteria for genetic studies of Tourette syndrome. *Archives of neurology*, *54*(5), 517-518.

Lakens, D. (2013). Calculating and reporting effect sizes to facilitate cumulative science: a practical primer for t-tests and ANOVAs. *Frontiers in psychology, 4*.

Lang, W. (2003). Surface recordings of the Bereitschaftspotential in normals. In *The Bereitschaftspotential* (pp. 19-34). Springer US.

Lecas, J. C., Requin, J., Anger, C., & Vitton, N. (1986). Changes in neuronal activity of the monkey precentral cortex during preparation for movement. *Journal of Neurophysiology, 56*(6), 1680-1702.

Lefebvre, G., Lebreton, M., Meyniel, F., Bourgeois-Gironde, S., & Palminteri, S. (2016). Asymmetric reinforcement learning: computational and neural bases of positive life orientation. *bioRxiv*, 038778.

Libet, B. et al. (1983) Time of conscious intention to act in relation to onset of cerebral activity (readiness-potential). The unconscious initiation of a freely voluntary act. *Brain 106,* 623–642.

Limousin, P., Greene, J., Pollak, P., Rothwell, J., Benabid, A. L., & Frackowiak, R. (1997). Changes in cerebral activity pattern due to subthalamic nucleus or internal pallidum stimulation in Parkinson's disease. *Annals of neurology, 42*(3), 283-291.

Lipsey, M.W., & Wilson, D. B. (2001). *Practical meta-analysis (Vol. 49).* Thousand Oaks, CA: Sage publications.

Lotze, M., Montoya, P., Erb, M., Hülsmann, E., Flor, H., Klose, U., ... & Grodd, W. (1999). Activation of cortical and cerebellar motor areas during executed and imagined hand movements: an fMRI study. *Journal of cognitive neuroscience, 11*(5), 491-501.

Louis, E. D. (2000). Essential tremor. *Archives of neurology, 57*(10), 1522-1524.

Loukas, C., & Brown, P. (2004). Online prediction of self-paced hand-movements from subthalamic activity using neural networks in Parkinson's disease. *Journal of neuroscience methods, 137*(2), 193-205.

Maeda, T., Kato, M., Muramatsu, T., Iwashita, S., Mimura, M., & Kashima, H. (2012). Aberrant sense of agency in patients with schizophrenia: forward and backward over-attribution of temporal causality during intentional action. *Psychiatry research, 198*(1), 1-6.

Manohar, S. G., Chong, T. T. J., Apps, M. A., Batla, A., Stamelou, M., Jarman, P. R., ... & Husain, M. (2015). Reward pays the cost of noise reduction in motor and cognitive control. *Current Biology, 25*(13), 1707-1716.

Maris, E., & Oostenveld, R. (2007). Nonparametric statistical testing of EEG-and MEG-data. Journal of neuroscience methods, 164(1), 177-190. Martinez-Martin, P., Rodriguez-Blazquez, C., Kurtis, M. M., & Chaudhuri, K. R. (2011). The impact of non-motor symptoms on health-related quality of life of patients with Parkinson's disease. *Movement Disorders, 26*(3), 399-406.

McClure, S. M., Berns, G. S., & Montague, P. R. (2003). Temporal prediction errors in a passive learning task activate human striatum. *Neuron, 38*(2), 339-346.

Meck, W. H. (2005). Neuropsychology of timing and time perception. *Brain and cognition, 58*(1), 1-8.

Metcalfe, J., & Greene, M. J. (2007). Metacognition of agency. *Journal of Experimental Psychology: General, 136*(2), 184.

Misirlisoy, E., & Haggard, P. (2014). Veto and vacillation: a neural precursor of the decision to withhold action. *Journal of cognitive neuroscience, 26*(2), 296-304.

Moore, J. W., & Haggard, P. (2008). Awareness of action: Inference and prediction. *Consciousness and cognition*, *17*(1), 136-144.

Moore, J. W., Lagnado, D., Deal, D. C., & Haggard, P. (2009). Feelings of control: contingency determines experience of action. *Cognition, 110*(2), 279-283.

Moore, J. W., & Obhi, S. S. (2012). Intentional binding and the sense of agency: a review. *Consciousness and cognition, 21*(1), 546-561.

Moore, J. W., Ruge, D., Wenke, D., Rothwell, J., & Haggard, P. (2010). Disrupting the experience of control in the human brain: pre-supplementary motor area contributes to the sense of agency. *Proceedings of the Royal Society of London B: Biological Sciences*, rspb20100404

Moore, J. W., Schneider, S. A., Schwingenschuh, P., Moretto, G., Bhatia, K. P., & Haggard, P. (2010). Dopaminergic medication boosts action–effect binding in Parkinson's disease. *Neuropsychologia, 48*(4), 1125-1132.

Moretto, G., Schwingenschuh, P., Katschnig, P., Bhatia, K. P., & Haggard, P. (2011). Delayed experience of volition in Gilles de la Tourette syndrome. *Journal of Neurology, Neurosurgery & Psychiatry*.

Neshige, R., Lüders, H., & Shibasaki, H. (1988). Recording of movement-related potentials from scalp and cortex in man. *Brain: a journal of neurology, 111*, 719-736.

Niv, Y., Edlund, J. A., Dayan, P., & O'Doherty, J. P. (2012). Neural prediction errors reveal a risk-sensitive reinforcement-learning process in the human brain. *The Journal of Neuroscience*, *32*(2), 551-562.

O'Doherty, J. P. (2014). The problem with value. *Neuroscience & Biobehavioral Reviews*, *43*, 259-268.

Ogloff, J. R. (1991). A comparison of insanity defense standards on juror decision making. *Law and Human Behavior, 15*(5), 509.

Olds, J., & Milner, P. (1954). Positive reinforcement produced by electrical stimulation of septal area and other regions of rat brain. *Journal of comparative and physiological psychology, 47*(6), 419.

Pastor, M. A., Artieda, J., Jahanshahi, M., & Obeso, J. A. (1992). Time estimation and reproduction is abnormal in Parkinson's disease. Brain, 115(1), 211-225.

Patzelt, E. H., Kurth-Nelson, Z., Lim, K. O., & MacDonald, A. W. (2014). Excessive state switching underlies reversal learning deficits in cocaine users. *Drug and alcohol dependence, 134,* 211-217.

Perneger, T. V. (1998). What's wrong with Bonferroni adjustments. *BMJ: British Medical Journal, 316*(7139), 1236.

Robertson, M. M., & Cavanna, A. E. (2009). The neuropsychiatry and neuropsychology of Gilles de la Tourette syndrome. *Neuropsychological Assessment of Neuropsychiatric Disorders.* 3rd ed. Oxford University Press, New York, 241-266.

Rolls, E. T. (2000). The Orbitofrontal Cortex and Reward. *Cerebral Cortex*, *10*(3), 284–294.

Roskies, A. L. (2010). How does neuroscience affect our conception of volition? *Annual review of neuroscience*, *33*, 109-130.

van der Salm, S. M., Tijssen, M. A., Koelman, J. H., & van Rootselaar, A. F. (2012). The bereitschaftspotential in jerky movement disorders. *Journal of Neurology, Neurosurgery & Psychiatry, 83*(12), 1162-1167.

Sankoh, A. J., Huque, M. F., & Dubey, S. D. (1997). Some comments on frequently used multiple endpoint adjustment methods in clinical trials. *Statistics in medicine, 16*(22), 2529-2542.

Santaniello, S., McCarthy, M. M., Montgomery, E. B., Gale, J. T., Kopell, N., & Sarma, S. V. (2015). Therapeutic mechanisms of high-frequency stimulation in Parkinson's disease and neural restoration via loop-based reinforcement. *Proceedings of the National Academy of Sciences, 112*(6), E586-E595.

Schall, J. D. (2002). The neural selection and control of saccades by the frontal eye field. *Philosophical Transactions of the Royal Society of London B: Biological Sciences, 357*(1424), 1073-1082.

Schultz, W. (2013). Updating dopamine reward signals. *Current opinion in neurobiology, 23*(2), 229-238.

Schüür, F., & Haggard, P. (2011). What are self-generated actions? *Consciousness and Cognition, 20*(4), 1697-1704.

Sescousse, G., Caldú, X., Segura, B., & Dreher, J. C. (2013). Processing of primary and secondary rewards: a quantitative meta-analysis and review of human functional neuroimaging studies. *Neuroscience & Biobehavioral Reviews, 37*(4), 681-696.

Shadlen, M. N. & Gold (2004), J. I. in *The Cognitive Neurosciences* 3rd edn (ed. Gazzaniga, M. S.) 1229–1241 MIT Press.

Shibasaki, H., Barrett, G., Halliday, E., & Halliday, A. M. (1980). Components of the movement-related cortical potential and their scalp topography. *Electroencephalography and clinical neurophysiology, 49*(3), 213-226.

Shibasaki, H. & Hallett, M. (2006) What is the Bereitschaftspotential? *Clinical. Neurophysiology. 117*, 2341–2356.

Sidarus, N., Chambon, V., & Haggard, P. (2013). Priming of actions increases sense of control over unexpected outcomes. *Consciousness and cognition, 22*(4), 1403-1411.

Simonetta, M., Clanet, M., & Rascol, O. (1991). Bereitschaftspotential in a simple movement or in a motor sequence starting with the same simple movement. *Electroencephalography and Clinical Neurophysiology/Evoked Potentials Section, 81*(2), 129-134.

Sirigu, A., Daprati, E., Pradat-Diehl, P., Franck, N., & Jeannerod, M. (1999). Perception of self-generated movement following left parietal lesion. *Brain, 122*(10), 1867-1874.

Skinner, B. F. (1938). *The behaviour of organisms: An experimental analysis.* New York. D. Appleton-Century Company Incorporated.

Smith, J. & Alloy, L., (2009) A roadmap to rumination: a review of the definition, assessment, and conceptualization of this multifaceted construct. *Clinical Psychology Review, 29(*2).

Smittenaar, P., Prichard, G., FitzGerald, T. H., Diedrichsen, J., & Dolan, R. J. (2014). Transcranial direct current stimulation of right dorsolateral prefrontal cortex does not affect model-based or model-free reinforcement learning in humans. *PloS one, 9*(1), e86850.

Spengler, S., von Cramon, D. Y., & Brass, M. (2009). Control of shared representations relies on key processes involved in mental state attribution. *Human brain mapping, 30*(11), 3704-3718.

Sperduti, M., Delaveau, P., Fossati, P., & Nadel, J. (2011). Different brain structures related to self-and external-agency attribution: a brief review and meta-analysis. *Brain Structure and Function, 216*(2), 151-157.

Starkstein SE, Mayberg HS, Preziosi TJ, Andrezejewski P, Leiguarda R, Robinson RG. (1992) Reliability, validity, and clinical correlates of apathy in Parkinson's disease. *Journal of Neuropsychiatry and Clinical Neurosciences*. *4*(2):134–139

Stevens, J. C., & Mack, J. D. (1959). Scales of apparent force. *Journal of experimental psychology, 58*(5), 405.

Sutton, R. S., & Barto, A. G. (1991). *Introduction to Reinforcement Learning.* MIT Press.

Swainson, R., Rogers, R. D., Sahakian, B. J., Summers, B. A., Polkey, C. E., & Robbins, T. W. (2000). Probabilistic learning and reversal deficits in patients with Parkinson's disease or frontal or temporal lobe lesions: possible adverse effects of dopaminergic medication. *Neuropsychologia, 38*(5), 596-612.

Synofzik, M., Vosgerau, G., & Newen, A. (2008). Beyond the comparator model: a multifactorial two-step account of agency. *Consciousness and cognition, 17*(1), 219-239.

Tabu, H., Aso, T., Matsuhashi, M., Ueki, Y., Takahashi, R., Fukuyama, H., Shibasaki, H. & Mima, T. (2015) Parkinson's disease patients showed delayed awareness of motor intention. *Neuroscience Research, 95*, 74-77.

Takahata, K., Takahashi, H., Maeda, T., Umeda, S., Suhara, T., Mimura, M., & Kato, M. (2012). It's not my fault: postdictive modulation of intentional binding by monetary gains and losses. *PloS one, 7*(12), e53421.

Tarkka, I. M., & Hallett, M. (1991). Topography of scalp-recorded motor potentials in human finger movements. *Journal of Clinical Neurophysiology, 8*(3), 331-341.

Toma K, Matsuoka T, Immisch I, Mima T, Waldvogel D, Koshy B, et al. (2002) Generators of movement-related cortical potentials: fMRI-constrained EEG dipole source analysis. *NeuroImage, 17,* 161–73.

Tomlinson, C. L., Stowe, R., Patel, S., Rick, C., Gray, R., & Clarke, C. E. (2010). Systematic review of levodopa dose equivalency reporting in Parkinson's disease. *Movement disorders, 25*(15), 2649-2653.

Turton, A. J., & Fraser, C. M. (1986). A test battery to measure the recovery of voluntary movement control following stroke. *International rehabilitation medicine, 8*(2), 74-78.

Uitenbroek, D. G. (1997). *SISA Bonferroni*. Southampton: D.G. Uitenbroek. Retrieved September 9, 2017, from the World Wide Web: http://www.quantitativeskills.com/sisa/calculations/bonfer.htm

Vaillancourt, D. E., Schonfeld, D., Kwak, Y., Bohnen, N. I., & Seidler, R. (2013). Dopamine overdose hypothesis: evidence and clinical implications. *Movement Disorders, 28*(14), 1920-1929.

Verbaarschot, C., Farquhar, J., & Haselager, P. (2015). Lost in time...: The search for intentions and readiness potentials. *Consciousness and cognition, 33*, 300-315.

Volkmann, J. (2004). Deep brain stimulation for the treatment of Parkinson's disease. *Journal of clinical neurophysiology, 21*(1), 6-17.

Vulkan, N. (2000). An economist's perspective on probability matching. *Journal of economic surveys, 14*(1), 101-118.

Walsh, E., & Haggard, P. (2013). Action, prediction, and temporal awareness. *Acta Psychologica*, *142*(2), 220–229.

Wegner, D. M., & Wheatley, T. (1999). Apparent mental causation: Sources of the experience of will. *American psychologist, 54*(7), 480.

Weintraub, D., David, A.S., Evans, A.H., Grant, J.E., Stacy, M. (2015) Clinical spectrum of impulse control disorders in Parkinson's disease, Journal of *movement disorders, 30* (2).

Weintraub, D., Mamikonyan, E., Papay, K., Shea, J. A., Xie, S. X., & Siderowf, A. (2012). Questionnaire for impulsive‐compulsive disorders in Parkinson's Disease‐Rating Scale. *Movement Disorders, 27*(2), 242-247.

Wenke, D., Fleming, S. M., & Haggard, P. (2010). Subliminal priming of actions influences sense of control over effects of action. *Cognition, 115*(1), 26–38.

Wolpe, N., Moore, J. W., Rae, C. L., Rittman, T., Altena, E., Haggard, P., & Rowe, J. B. (2014). The medial frontal-prefrontal network for altered awareness and control of action in corticobasal syndrome. *Brain, 137*(1), 208–220.

Wunderlich, K., Smittenaar, P., & Dolan, R. J. (2012). Dopamine enhances model-based over model-free choice behavior. *Neuron, 75*(3), 418-424.

Wundt, W. M. (1907). *Outlines of psychology*. W. Engelmann.

Yoshie, M., & Haggard, P. (2013). Negative emotional outcomes attenuate sense of agency over voluntary actions. *Current Biology, 23*(20), 2028-2032.