The role of dopamine in learning, movement & motivation

Tamar Shiner, MBChB MRCP

Institute of Neurology

University College London

Thesis submitted for the degree of Doctor of Philosophy (PhD) at University College London

Primary Supervisor: Professor Raymond J Dolan, FRS

Secondary Supervisor: Professor Kailash P Bhatia, MD DM FRCP

September 2011

I, Tamar Shiner, 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

Signature:

Abstract

The primary aim of the research I have undertaken is to better understand the influence of dopamine on behavior and to build on knowledge of the various roles of dopamine in the healthy brain but also to improve understanding of the deficits affecting patients with Parkinson's disease (PD), the hallmark of which is dopamine depletion.

By testing PD patients on cognitive and motor tasks, we are able to probe the effects of dopamine depletion in humans. Testing PD patients in different medication states also provides a method with which to attempt to tease apart the various roles of dopamine from each other. My first two experiments use the PD model to this end whereas the third experiment utilises a pharmacological manipulation in healthy individuals.

The aim of my first experiment was to tease apart the relative contribution of dopamine to learning from its influence on action performance, and by doing this to better understand the deficits which have been observed in PD patients in reinforcement learning tasks.

The second experiment focuses on the motor deficits observed in PD. The aim of this study was to test whether these motor deficits can at least in part explained by the deficits in reward sensitivity.

The third and final experiment in this thesis uses a pharmacological manipulation in healthy individuals to isolate the role of dopamine in set shifting in the context of a response to cues with negative hedonic valence, with the hope of better

understanding the neurobiology underlying pathological behaviours associated with the hyperdopaminergic state.

Acknowledgements

This research would not have been possible without the support of many people. I feel honored to have been in the position to work with such leaders in the field who took the time to teach and guide me through this work.
These are some, but not all of the people who have made this work possible.

My supervisor, Professor Raymond Dolan, for taking me on, helping me secure funding and supporting me throughout. His help at every stage of this research was invaluable.

My second supervisor, Professor Kailash Bhatia, who helped me with recruitment and gave me the opportunity to participate in his fascinating movement disorders clinics and for providing general inspiration.

Professor Peter Dayan, whose ideas and suggestions shaped much of this work and whose thoughtful input has been invaluable.

Professor Karl Friston, for the incredible advice and teaching and for the opportunity to absorb (as much as I was able) his inspirational and groundbreaking ideas.

My wonderful friends, Dr Tali Sharot and Dr Rosalyn Moran who apart from being the best friends one could hope for, are also brilliant neuroscientists who have taught me so much.

My other close colleagues and friends, Dr Marc Guitart-Masip, Dr Mkael Symmonds, Dr Nick Wright and Dr Steve Fleming who have provided immense help and also made working at the FIL so much fun.

Dr Ben Seymour for helping me write the grant proposal, which finally secured funding and then providing advice and assistance throughout my research.

Dr Klaus Wunderlich for his fantastic input, technical help and great suggestions.

My exceptionally supportive and loving parents, Claudie and Rob, who always believed in me and taught me to pursue my ambitions and to never give up.

My wonderful husband Jonathan who provided immense support, patience and love allowing me to pursue this project and of course my fantastic son Jake who helped with the typing and dutifully came everywhere with me while I was testing subjects for a whole 9 months!

Finally, the Medical Research Council who provided funding for this work.

Table of Contents

Abbreviations

AADI	Aromatic-amino-acid de-carboxylase inhibitor
ACC	Anterior cingulate cortex
ANOVA	Analysis of variance
BDI	Becks Depression Inventory
BOLD	Blood-oxygen-level dependence
COMT	Catechol-O-methyl transferase
DA	Dopamine
DAT	Dopamine active transporter
DRT	Dopamine replacement therapy
FDR	False discovery rate
fMRI	Functional magnetic resonance imaging
FWE	Family wise error
GABA	gamma-Aminobutyric acid
GLM	General linear model
GPe	External segment of the globus pallidus
GPi	Internal segment of the globus pallidus
HRF	Haemodynamic response function
ICD	Impulse control disorder
ITI	Inter-trial interval
L-dopa	Levodopa
LTD	Long term depression
LTP	Long term potentiation

MAO	Monoamine oxidase
MMSE	Mini Mental State Examination
MNI	Montreal neurological institute
MRI	Magnetic resonance imaging
MT	Movement time
NAc	Nucleus accumbens
NMR	Nuclear magnetic resonance
OFC	Orbitofrontal cortex
PD	Parkinson's disease
PET	Positron emission tomography
RL	Reinforcement learning
ROI	Region of Interest
RRF	Retrorubral field
RT	Reaction time
SD	Standard deviation
SE	Standard error
SNc	Substantia nigra pars compacta
SNr	Substantia nigra pars reticulate
SPM	Statistical parametric mapping
STN	Subthalamic nucleus
vmPFC	Ventromedial prefrontal cortex
VP	Ventral pallidum
VTA	Ventral tegmental area

Chapter 1

Introduction

1.1 Outline of thesis

Dopamine is a central neurotransmitter in the basal ganglia and influences many different aspects of behaviour. The primary aims of the research I have undertaken are to: (i) build on the knowledge of role of dopamine in the healthy brain and specifically its influence on cognition; (ii) improve our understanding of the cognitive deficits affecting patients with Parkinson's disease (PD).

As PD is characterised by dopaminergic cell loss, it provides a valuable human model of the dopamine depleted state. By observing deficits present in PD patients and the effect of dopaminergic medication on these deficits, it is possible to make inferences about the function of dopamine in the healthy brain while also learning about the manifestations of a hypodopaminergic state in this patient population.

PD is a common, progressive degenerative neurological disorder (Pavese & Brooks 2009) with prevalence rates standing at 1.8% in people over the age of 65 (de Rijk et al 2000). The central feature of PD is loss of dopaminergic pigmented neurons in the substantia nigra pars compacta (SNc) leading, among other effects, to decreased levels of dopamine in the striatum (Koller & Melamed 2007b). By testing PD patients on cognitive and motor tasks, I was able to probe the effects of dopamine depletion in humans. Testing PD patients ON and OFF dopamine replacement therapy (DRT) provides a method with which to attempt to tease apart the various roles of dopamine from one another, and ultimately to answer key questions surrounding how dopamine modulation exerts its effects on aspects of cognition, movement, and the interplay between the two.

The first two experiments in this thesis use the PD model to this end while the third experiment utilises a pharmacological manipulation in healthy individuals.

The aim of my first experiment was to distinguish the influence of dopamine on learning from its influence on action performance and thus better understand the relative contribution of dopamine to learning and performance. In doing so, we hoped to explain the deficits which have been observed in PD patients, as seen for example in reinforcement learning tasks.

The second experiment examines the motor deficits observed in PD. The aim of this study was to test whether the motor deficits found in PD patients can, at least in part, be explained by insensitivity to reward observed in this group, suggesting that there may be a carryover from the cognitive to the motor domain in PD.

In order to probe the effect of dopaminergic modulation further, the third and final experiment in this thesis uses a pharmacological manipulation in healthy individuals to isolate the role of dopamine in set shifting in the context of a response to cues with negative hedonic valence, with the hope of better understanding the neurobiology underlying pathological behaviours associated with the hyperdopaminergic state.

It is impossible to investigate the role of dopamine without understanding the basic neuroanatomy and physiology of the basal ganglia. I begin this thesis by considering this neuroanatomical background and giving an overview of the biochemistry of dopamine synthesis. I then discuss the seminal research in the area and conclude this section with a discussion on PD.

1.2 Neuroanatomy of the basal ganglia and dopamine system

1.2.1 Overview

The basal ganglia are a group of subcortical nuclei which have long been linked to movement control. As such, diseases affecting the basal ganglia (such as PD) are primarily associated with movement disorders (Koller & Melamed 2007b), although it is now clear that these diseases are also associated with cognitive deficits (Lees & Smith 1983). The basal ganglia are often divided into the dorsal division which consists of the neostriatum (caudate nucleus and putamen), the external and internal segments of the globus pallidus (GPe and GPi), the subthalamic nucleus (STN), and the substantia nigra pars reticulate (SNr) and pars compacta (SNc) and a ventral division consisting of nucleus accumbens (NAc), the ventral pallidum (VP) and ventral tegmental area (VTA) (Koller & Melamed 2007b).

Each basal ganglia nucleus is histologically distinct. The most abundant striatal cell type is the GABAergic medium spiny projection neuron, which represent 90–95% of all striatal neurons. These cells receive inputs from the cortex and thalamus and also intrinsic connections from inhibitory striatal interneurons, including cholinergic neurons. In addition they receive modulatory inputs from the dopaminergic and sertoninergic cells originating in the midbrain. The dopaminergic projection fibres terminate on the neck of the dendritic spines of striatal medium spiny neurons and are thus in a position to modulate corticostriatal information flow. As opposed to the heterogenous histological organisation of the striatum, The GPe and GPi are more homogeneous and contain only very few interneurons. The STN is distinct from the other basal ganglia nuclei due to its densely packed structure whose

neurons are excitatory and glutamatergic (Koller & Melamed 2007b). As well as dopamine cells (Swanson 1982) all divisions of the midbrain also contain substantial amounts of GABAergic cells (Nair-Roberts et al 2008). The VTA receives excitatory inputs from widely distributed brain areas indicating that rather than the VTA being influenced by a discrete set of brain structures it is likely to be regulated by an integrated network of inputs (Geisler & Zahm 2005). The only cortical structure however with a major projection to the VTA originates from the PFC (Sesack & Grace 2010) however the exact function of this pathway remains unclear although it appears to play a role in plasticity of dopamine neurons (Wolf et al 2004). It has also been demonstrated that PFC axons synapse onto dopamine neurons that then project back to the PFC creating a circuit which allows the PFC to regulate the extent of its modulatory feedback by dopamine (Carr & Sesack 2000). There is also major inhibitory feedback from the basal ganglia to the VTA and SNc (Sesack & Grace 2010). Recently a major ascending source of inhibition from the mesopontine rostromedial tegmental nucleus to the SNc has been discovered. This structure, which consists primarily of GABA cells, receives afferents from many forebrain and brainstem structures and has widespread projections to the SNc and VTA and is therefore in a critical position to inhibit dopamine cell firing in response to aversive stimuli (Jhou et al 2009). In addition, serotonin neurons in the dorsal raphe nucleus synapse onto dopaminergic cells (Van Bockstaele et al 1994). In summary, the VTA receives influences from multiple ascending, descending and intrinsic sources however the functional significance of each afferent in relation to reward has yet to be fully determined (Sesack & Grace 2010).

1.2.2 Loops of the basal ganglia – extrinsic connections

The basal ganglia receive glutametergic inputs from the cortex and the thalamus and have numerous intrinsic connections. The main input structure of the basal ganglia is the striatum and the main output structures are the GPi and SNr (Koller & Melamed 2007b). The relaying of information through the basal ganglia is implemented through spatially segregated but partially overlapping 'loops' from the cortex to the basal ganglia, through to the thalamus and back to the cortex (Alexander et al 1986). The overlapping of connections from cortical sites to different basal ganglia subregions is thought to allow for both parallel and integrative networks (Draganski et al 2008; Haber 2003), as well as for the channeling of information across these functional circuits (Haber 2003). Information is also channelled between the separate basal ganglia loops at the relays of the basal ganglia and the thalamus (McFarland & Haber

2002; Redgrave et al 2010).

There is a strong topographical organisation of the projections from the cortex to the striatum. In primates, the motor and somatosensory cortices project to the postcommissural putamen, the associative cortices to the caudate nucleus and the precommissural putamen. The limbic cortices, the amygdala and the hippocampus terminate preferentially in the ventral striatum (see figure I.I for schematic depiction).

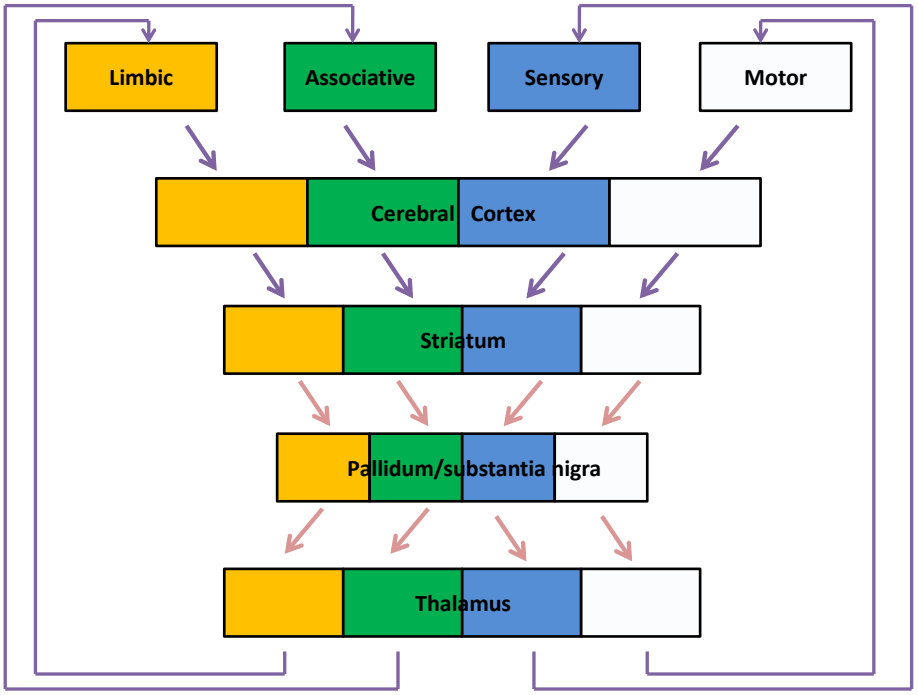

***Figure I.I**: Represented in this figure are parallel projecting basal ganglia loops which convey information from the cortex via the basal ganglia and thalamus conveying limbic (red/yellow), associative (yellow/green) and sensorimotor (green/blue) information. The topography of the territories are maintained throughout the basal ganglia and thalamic relays, albeit with overlap of connections from the cortex. There is an opportunity for modulation of activity inside these loops specifically at relays in the thalamus and basal ganglia and also by virtue of partially overlapping cortical connections. The purple arrows indicate glutametergic projections while pink arrows represent GABA-ergic projections. (Figure reproduced from Redgrave et al (Redgrave et al 2010)).*

As well as the corticostriatal projections there are also thalamostriatal projections which provide a major source of excitatory afferents. In addition there is a fast (single-synaptic) connection between the cortex and thalamus through the subthalamic nucleus that is thought to be excitatory.

Each spiny neuron cell receives synapses from thousands of distinct cortical neurons. This anatomical organisation is consistent with the idea that spiny cells integrate information from many sources (Kincaid et al 1998). Activity in striatal neurons is sensitive to context. This property has been demonstrated by experiments where striatal neurons fire in conjunction with a specific movement made in one context (e.g. self initiated) but are silent when an identical movement is made in another context (e.g. sensorally or memory guided) (Kimura et al 1992).

The suggestion that the NAc acts as a limbic-motor interface was first introduced by Mogenson et al (Mogenson et al 1980). This has been supported by evidence that the NAc receives inputs from limbic structures such as the amygdala, hippocampus and prefrontal cortex and sends outputs to output nuclei considered primarily motor (Nicola 2007), as well as evidence that dopamine injections into the NAc increase locomotion (Wu & Brudzynski 1995). The concept of this interface and a consequential suggestion that limbic inputs exert an influence on motor output is a central concept to this thesis.

1.2.3 Direct and indirect pathways in the basal ganglia – intrinsic connections

As well as extrinsic loops linking the basal ganglia to the cortex and thalamus, information is also conveyed through intrinsic loops within the basal ganglia. These intrinsic loops also maintain a high degree of spatial topography (Redgrave et al

2010; Wiesendanger et al 2004). Projections from the striatum to the globus pallidus and the substantia nigra maintain a general arrangement throughout these intrinsic nuclei. The posterior putamen is primarily engaged in sensorimotor functions, the caudate and anterior putamen primarily in associative functions and the ventral (limbic) striatum in motivational and emotional functions.

Within each of these circuits there are two central pathways, referred to as the direct and indirect pathways. The GABAergic medium spiny neurons of the *direct* pathway project to the neurons of the GPi and SNr, express substance P and dynorphin and preferentially express dopamine D1 receptors. The GABAergic medium spiny neurons of the *indirect* pathway project to the GPe, either directly or via the intercalated STN to GPi and SNr (Factor & Weiner 2008). The neurons of the *indirect* pathway express enkephalin and dopamine D2 receptors (figure I.II.A). Emerging evidence suggests that this model, whereby there is a clear differentiation between the direct and indirect pathways, as proposed by Albin et al in 1989 (Albin et al 1989) may be too simplistic. There is now persuasive evidence in support of collateral interactions across the system (Matamales et al 2009; Redgrave et al 2010; Smith et al 1998) and a subgroup of medium spiny neurons which co-express D1 and D2 receptors (Aizman et al 2000) (figure I.II.B for expanded model of function).

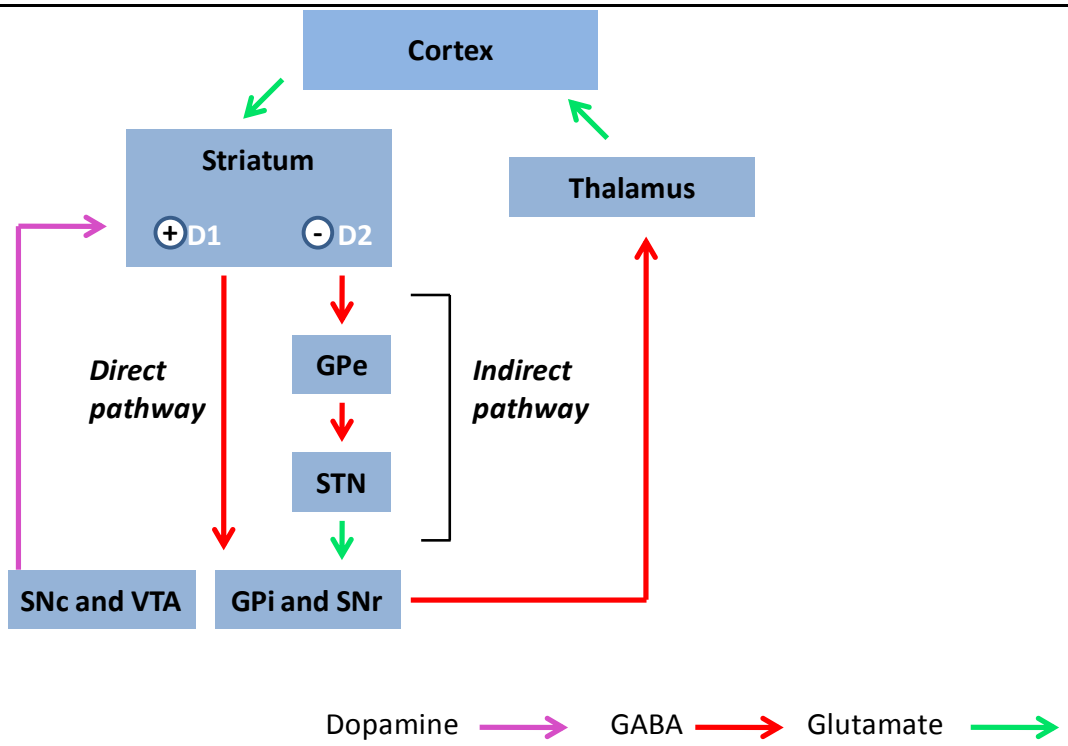

Figure I.II.A: *Original model of the direct and indirect pathways proposed by Albin et al (Albin et al 1989). In this model the output of the basal ganglia is though to be modulated by a balance (mediated by dopaminergic inputs from the SNc and VTA) between the direct and indirect pathways, which express primarily D1 or D2 receptors respectively. The direct pathway promoting behaviour and the indirect pathway suppressing behaviour via relays in the GPe and STN.*

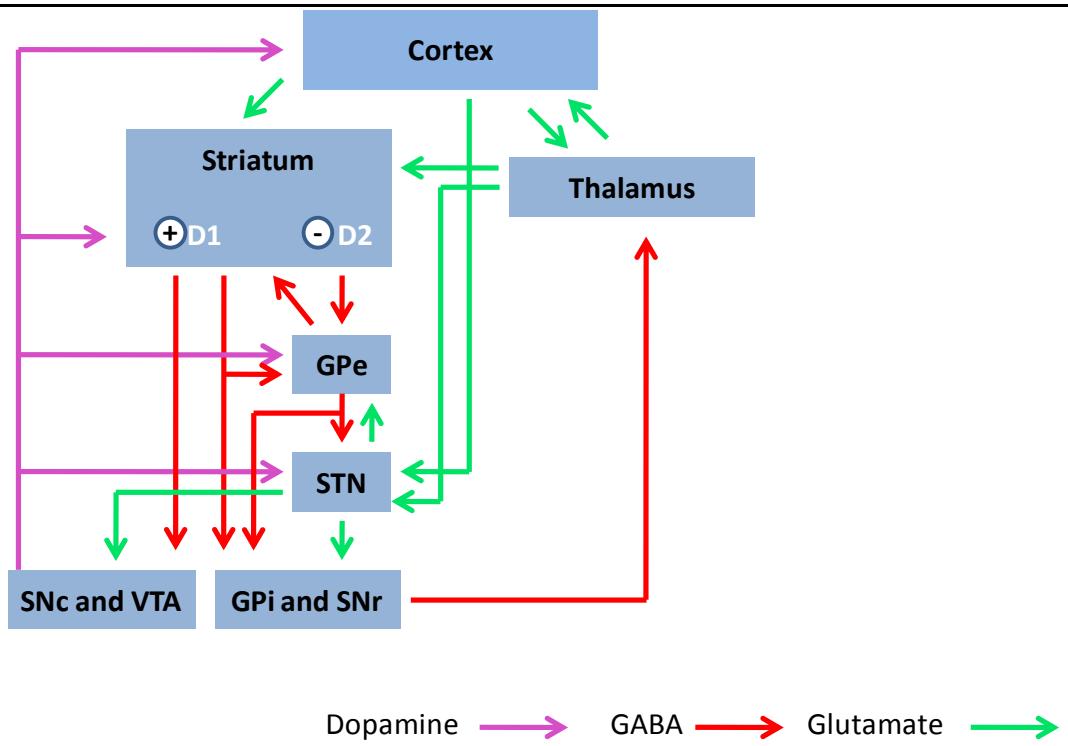

Figure I.II.B: *More recent anatomical investigations of the basal ganglia reveal many more connections which make the responsiveness of the system to inputs more difficult to directly predict. (Diagram reproduced from (Redgrave et al 2010)*

1.2.4 Dopaminergic projections

The SNc, VTA and the retrorubral field (RRF) provide the main sources of dopamine to basal ganglia. The SNc and RRF project to the caudate and putamen and form the nigrostriatal system. Neurons arising form the VTA form the mesolimbic and mesocortical pathways. In the mesolimbic system, neurons from the VTA project to the nucleus accumbens, olfactory tubercle, hippocampus, amygdala and septum whereas in the mesocortical pathway, VTA projects to cortical structures including the prefrontal, cingulate and perirhinal cortices (Arias-Carrion & Poppel 2007) (figure I.III). As a result of overlap of connections between the ventral striatum and the various tiers of the midbrain, the ventral striatum is in a position to influence activity in the more dorsal striatal regions, allowing for a limbic influence on motor regions (Haber 2003; Haber & Knutson 2010).

Dopaminergic inputs functionally regulate the activity of the striatal medium projection neurons via interactions with dopamine receptors (see below for more detailed discussion of dopamine receptors), and thereby exert an effect on corticostriatal transmission. The interaction between the glutamatergic cortical inputs and the modulatory dopaminergic projections occurs at postsynaptic dendritic spines (Smith & Bolam 1990). The glutametergic terminals converge on the head of the dendritic spines whereas the dopamergic terminals converge on the neck of the dendritic spines, providing a mechanism whereby dopaminergic inputs filter the more distal glutametergic afferents (Smith & Bolam 1990).

Copyrighted Figure removed

Figure I.III: *Dopamine projections to the forebrain. Illustrated are projections from the ventral tegmental area (VTA) to the nucleus accumbens and prefrontal cortex, and projections from the substantia nigra (SNc) to the dorsal striatum (caudate nucleus and putamen and related structures) extracted from (Arias-Carrion & Poppel 2007)*

1.3 Biochemistry of dopamine synthesis and action

1.3.1 Dopamine synthesis

Dopamine is a catecholamine with its own unique function within the mammalian brain. It is also a precursor for noradrenaline and adrenaline (figure I.IV), the other two naturally occurring brain catecholamines.

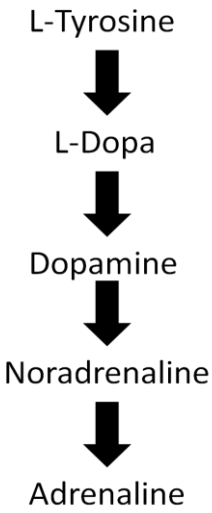

Figure I.IV: *Biosynthetic pathway for catecholamines. Reproduced from (Albers et al 2006).*

Once synthesised, dopamine is concentrated in storage vesicles, a high density of which are in nerve terminals (Albers et al 2006). The vesicles are responsible for maintaining a steady supply of synaptic catecholamines, by fusing with the membrane of the nerve terminal following the influx of calcium caused by an action potential

There are several factors which determine the duration of dopamine's action in the synapse. The first is the *amount* of dopamine released, the second is active *removal* of dopamine from the synaptic cleft by dopamine transporters (DAT), and a third is activation of *presynaptic receptors*, known as autoreceptors, which monitor the amount of dopamine in the synaptic cleft. When the autoreceptors are activated they inhibit synaptic release of dopamine (Koller & Melamed 2007b).

The metabolism of dopamine is effected by two pathways, both of which involve Monoamine oxidase (MAO) and catechol-O-methyltransferase (COMT) (Koller & Melamed 2007b). Both MAO and COMT act intercelluarly, inactivating the catecholamines which are not stored within vesicles, i.e. that are free within the nerve terminals. COMT is an enzyme which methylates catecholamines. Its expression varies among the normal population due to a common polymorphism in which valine is substituted for methionine. The valine polymorphism causes increased enzymatic activity and is associated with lower levels of prefrontal dopamine, impaired prefrontal cortical function and a reported increase in risk of developing schizophrenia (Albers et al 2006).

1.3.2 Dopaminergic transmission in the basal ganglia

The physiological actions of DA are mediated by at least five distinct G protein coupled receptor subtypes (Missale et al 1998). These receptors fall into two classes, the D1-like receptor subtypes (D1 and D5) which couple to the G protein and activate adenylyl cyclase, and the D2-like receptor subfamily that inhibit adenylyl cyclase and activate K^+ channels (Berke & Hyman 2000; Ebadi & Pfeiffer 2005). In the striatum there is a very high concentration of D1 and D2 dopamine

receptors, localised concentrations of D3 receptors in the ventral striatum and lower levels of the D4 and D5 type receptors (Bloom & Kupfer 1995). Unlike the excitatory effect of D1 receptor activation, D2 stimulation activates inwardly rectifying K+ channels and leads to a *decrease* in cell excitability (Berke & Hyman 2000; Koller & Melamed 2007b).

Striatal D2 receptors are tonically (continuously) stimulated by basal levels of dopamine. This tonic activity is important for normal motor behaviour. Accordingly, mice lacking D2 receptors are parkinsonian (Baik et al 1995) in a similar way to animals that are given D2 antagonists. Of note, mice bred to lack D1 receptor expression do not show parkinsonian symptoms (Drago et al 1994).

A further important factor is the affinity of the various dopamine receptors for dopamine, their substrate. The D1 family of receptors have a low affinity for dopamine whereas the D2 like receptors have a high affinity (Creese et al 1983). As a result of these differences in affinities it is thought that a high level of phasic dopamine release is required for D1 receptor activation, whereas the D2 receptors, which have a high affinity for dopamine, are continuously activated by the lower, tonic levels of dopamine (Grace 1991).

Dopamine neurons in vivo have 3 patterns of activity, an inactive hyperpolarised state, a tonic (single spike) or a phasic (burst) mode (Grace et al 2007). The tonic firing is controlled by an intrinsic pacemaker whereas the burst firing is reliant on afferent input. Studies suggest that the neurons of the VTA are held in an inactive hyperpolarised state by the ventral pallidum. The ventral pallidum has a high rate of spontaneous activity which inhibits VTA dopamine neurons. The release of this

inhibition leads to spontaneous firing that provides the baseline levels of extrasynaptic dopamine referred to as the tonic level of dopamine (Grace et al 2007). The burst firing rate is dependent on afferent inputs and this fact has lent weight to theories implicating it in reward related and goal directed behaviours (Grace et al 2007).

In addition, D1 receptors have an important role in plasticity whereby activation of these receptors leads to long-term potentiation (LTP) and positive reinforcement of behaviour (Reynolds et al 2001). It has been suggested that the positive phasic bursting of dopamine neurons activates D1 receptors and induces LTP and long term depression (LTD) via D2 receptors, with the reverse occurring when there is a pause in dopamine neuron firing (Maia & Frank 2011).

The D3 receptors are expressed more densely in the limbic striatum then in other brain regions, with the NAc and associative striatum showing the highest concentration of D3 receptors of all brain regions (Levant 1997). Both dopamine agonists, primarily used to treat the motor effects of PD, and the antipsychotics, primarily used to treat the affective disorders such as schizophrenia, usually have dual action on both the D2 and D3 receptors, stimulating and blocking them respectively (Sokoloff et al 1990). Dopamine agonists with high D3 receptor affinity have an association with cognitive side effects such as pathological gambling (Dodd et al 2005), however dopamine agonists with different receptor affinity profiles can also cause cognitive side effects thereby complicating this direct association (Koller & Melamed 2007b).

1.4 Roles of dopamine

As detailed above, dopaminergic inputs modulate corticostriatal transmission and as such dopamine is a crucial neurotransmitter in the basal ganglia. Dopamine has been identified as being involved in many different and partially overlapping aspects of behaviour, although the integration of these aspects with one another continues to be a subject for debate. One of the central debates surrounds whether dopamine cell activity is (i) the driving force behind reward learning; (ii) the consequence of learning with its main role being on the expression of learning via performance effects, such as modulation of motivated behaviour; or (iii) whether it is on a combination of these functions.

The following is a brief summary of the key behaviours with which dopamine is thought to be involved: learning; invigoration of motivated behaviour; movement and action initiation.

1.4.1 Learning

Dopamine neurons of the VTA and SNc are involved with the processing of rewarding stimuli. Dopamine neurons respond with short, phasic activations when monkeys are presented with appetitive stimuli. However when these rewards become predicted by a cue, dopamine neurons change the time of their phasic activation from time of reward delivery to time of cue onset (Hollerman & Schultz 1998; Schultz et al 1997). If the predicted reward is then not delivered, then dopamine neuron activity becomes depressed below the basal firing rate at exactly the time that the reward should have been delivered. These findings, from the

seminal work of Schultz et al (Hollerman & Schultz 1998; Schultz et al 1997), demonstrate that rather than simply signalling rewards, dopamine neurons code for a difference between the expected and delivered reward, in other words the 'prediction error' (Schultz et al 1997) (figure I.V).

Copyrighted Figure removed

Figure I.V: *Represented in this figure are the seminal neurophysiological results of Schultz et al. The top panel shows midbrain dopamine neuron activity before learning takes place when a monkey is given a drop of unexpected appetitive juice. At this point there is no prediction so the activity in the midbrain dopaminergic neurons report a positive prediction error when the reward is delivered. The middle panel shows dopamine neuron activity when learning has taken place. Now the dopamine cell firing occurs at the same time as the cue predicting the reward and therefore when reward delivery is fully predicted the dopamine neurons are not activated by the delivery of reward. The bottom panel shows dopamine neuron activity after learning has taken place. As in the middle panel, when the reward-predicting cue is presented, dopamine cell activity increases. When the predicted reward subsequently fails to be delivered, however, there is a dip in firing at exactly the time when the reward should have been delivered (negative prediction error) (taken from (Schultz et al 1997)).*

It has also been highlighted that the prediction error response resembles the 'teaching signal' derived from computational reinforcement learning theory (the Rescorla-Wagner learning model) (Moll et al 1999; Schultz 2010; Schultz et al 1997), providing further support for its role in guiding learning via synaptic plasticity (Centonze et al 2001; Maia & Frank 2011). Experimental evidence linking dopamine to reward learning is extensive (Bayer & Glimcher 2005; Hollerman & Schultz 1998; Schultz et al 1997; Wise 2004; Wise & Rompre 1989) and much work supporting a role for dopamine signalling error prediction has emerged (Montague et al 2004; Schultz et al 1997).

The behaviour of dopamine neurons in response to predicted and surprising rewards, as described above, has been widely replicated. The proof, however, that this is the critical signal which drives learning and causes "stamping in" of stimulus-cue and stimulus-response associations is more controversial (Berridge & Robinson 1998). Contributing to the controversy are studies in which dopamine deficient mice have been shown to have no deficits in learning in response to rewards and they are capable of picking the most rewarding drink (sucrose compared with water) when presented with both (Cannon & Palmiter 2003). Other evidence against dopamine being required for learning has come from work on genetically dopamine deficient mice which show that although learning in a maze task appeared initially impaired in this group, when they were subsequently treated with L-dopa (Robinson et al 2005), or caffeine (Hnasko et al 2005) these lesioned mice had learned, an effect one can frame as demonstrating an effect of dopamine on the expression of learning rather than on learning itself.

There have been several studies demonstrating a boosting of reward prediction errors and differences in learning rates in hypo-, hyper- and normal dopaminergic states (Pessiglione et al 2006; Voon et al 2010). This is in contrast to the results of experiments on hyperdopaminergic mutant mice who show no difference in the speed of action-outcome learning (Yin et al 2006), In one of these studies however, the hyperdopaminergic rats pressed the lever more frequently than the normal rats in order to obtain rewards, a fact attributed to dopamine having an effect on performance rather than learning (Yin et al 2006).

As a consequence of these, and other findings, it has been suggested that the firing of dopamine neurons, rather than acting as a teaching signal which enhances learning, may instead be an actual consequence of learning which takes place elsewhere in the brain. These theories suggest that instead, dopamine serves to attach 'incentive motivation' to learned cues (Berridge 2007), which has been defined as "a conditioned motivation response of a brain, usually triggered by and assigned to a reward-related stimulus"(Berridge 2007; Berridge & Robinson 1998).

1.4.2 Motivation and vigour

It has been proposed that tonic dopamine levels are involved in controlling movement rate and vigour (Salamone & Correa 2002; Ungerstedt 1971). Dopamine has been shown to be involved with motivational engagement and vigour (Bardgett et al 2009; Berridge & Robinson 1998; Boureau & Dayan 2011; Lex & Hauber 2010; McClure et al 2003b; Niv 2007; Niv et al 2007), with dopamine depletion causing decreased motivation to work for rewards when reinforcement schedules are

demanding (Niv 2007; Salamone & Correa 2002) and, conversely, invigorating actions when dopamine levels are high (Salamone et al 2005).

Dopamine levels have been shown to be involved in determining the amount of effort an animal is prepared to make to obtain a reward and to have an effect on the rate of responding (Niv 2007; Salamone et al 2003). These findings fit with theories derived from computational neuroscience which suggest that this pattern reflects the net rate of environmental rewards being signalled by tonic dopamine levels. Thus, in a high tonic dopamine state, a high net rate of environmental reward is signalled and consequently every second during which a reward is not reaped is more costly (a cost of sloth) making it worthwhile for subjects to perform actions quickly even if the energetic cost is higher (Niv 2007; Niv et al 2007). The incentive salience point of view fits with this account as according to this theory dopamine is involved with the attachment of motivational properties to rewarding cues (Berridge 2007).

1.4.3 Action selection and movement

Finally, it is clear that dopamine plays a key role in action selection and movement, although precisely how this function operates is still subject of debate.

A popular model for this is the actor-critic model (Joel et al 2002; O'Doherty et al 2004). In this model the cortex is thought to represent the current state and the basal ganglia implements two computational roles – the 'critic' (presumed to be implemented in the ventral striatum and possibly in the PFC and amygdala) which learns the values of outcomes and the 'actor' (presumed to be implemented in the dorsal striatum) which learns stimulus-response associations so that actions

associated with long term rewards are subsequently chosen more frequently, with both regions updating their estimates based on dopaminergic prediction errors.

A different view is that the basal ganglia arbitrate between actions that are under consideration by the cortex so as to facilitate the best one based on learnt reinforcement probabilities (learnt action values) (Samejima et al 2005).

Another important factor when considering the modulatory functions of dopamine is the relative influence of phasic versus tonic dopamine, with phasic dopamine neuron activity being associated with learning from rewards by enhancing Go activity via low affinity D1 receptors and conversely tonic dopamine dips driving NoGo behaviour via the high affinity D2 receptor stimulation when there are pauses in firing (Bayer & Glimcher 2005; Frank et al 2007; Frank et al 2004; Maia & Frank 2011). It has been proposed that they work in concert to obtain rewards, with D1 receptors being stimulated by phasic firing and, at the same time, D2 receptors being stimulated by tonic dopamine thereby reducing activity in the NoGo pathway which allows facilitation of reward directed behaviours (Hikida et al 2010).

What is clear from the above discussion is that no overall consensus has as yet emerged regarding the precise nature of the influence of dopamine on behaviour.

1.5 Serotonin and behaviour

Many neurotransmitters aside from dopamine are involved in cognition, however much research has focussed on a central role for serotonin, especially in the modulation of aversive learning and decision making. However, the way in which serotonin influences behaviour is complex. Part of this complexity arises from studies which have demonstrated that both medications which reduce serotonin transmission, such as benzodiazepines (Deakin & Graeff 1991), and the chronic use of medications which increase serotonin concentrations, such as selective serotonin reuptake inhibitors (Hollander 1998), have anxiolytic effects. Indeed both class of drugs are effective in treating disorders such as anxiety and depression which are in turn associated with increased aversive processing (Cools et al 2008a).

Furthermore, acute tryptophan depletion, a procedure which reduces levels of serotonin, has been shown to increase BOLD activity in response to punishment prediction but to have no effect on reward predictions (Cools et al 2008b). It has been suggested that serotonin might impact punishment prediction errors much in the way that dopamine has been proposed to impact positive prediction errors (Daw et al 2002).

As is the case with dopamine, serotonin has effects on behavioural inhibition and its depletion has been shown to worsen behavioural responses to stimuli which require motor inhibition. In addition, serotonin depletion has been shown to impair acquisition in a conditioned visual discrimination task (Harrison et al 1999), thereby having effects on both learning and performance.

The various roles of serotonin and their interactions with dopamine are still not fully understood. There remain several paradoxes such as serotonin depletion being implicated in diseases characterised by impulsivity and *reduced* aversive processing while also being central in diseases such as depression characterised by reduced behavioural vigour and *enhanced* aversive processing (Cools et al 2011).

1.6 Parkinson's disease

1.6.1 Neuroanatomy of Parkinson's disease

Parkinson's disease is a progressive degenerative neurological disorder characterised by asymmetric onset of tremor, rigidity and bradykinesia (Pavese & Brooks 2009). It is one of the most common neurological disorders with prevalence rates standing at 1.8% in people over the age of 65 (de Rijk et al 2000). The cardinal pathological feature of PD is loss of pigmented neurons in the substantia nigra pars compacta leading to decreased levels of dopamine in the striatum (Koller & Melamed 2007b). The SNc does not directly participate in the transfer of information along the basal ganglia thalamocortical pathways, but is part of the brainstem catecholaminergic systems, providing dopaminergic inputs to striatum and other targets (Koller & Melamed 2007b).

The pattern of degeneration in PD is uneven with the posterior putamen and ventrolateral substantia nigra being affected first (Cools 2006; Pavese & Brooks 2009). Fluorodopa PET imaging has revealed that the largest uptake reduction is in dorsal posterior putamen

contralateral to the side of maximal clinical symptom expression (Morrish et al 1995) (figure VI).

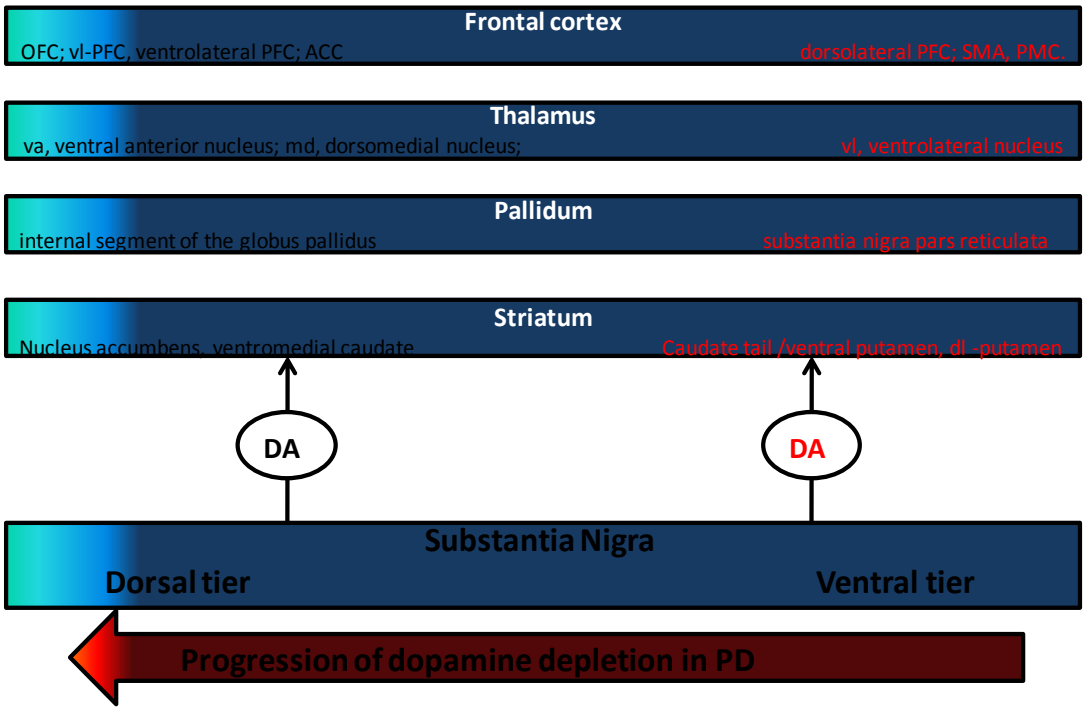

Figure VI: *Schematic to show the progession of dopamine cell degeneration in PD. In PD there is a spatio-temporal progression of dopamine (DA) cell degeneration from the ventral tier to the dorsal tier of the midbrain (which includes the VTA). The dark blue to light green shading gradient represents the spatio-temporal progression of pathology from dorsal to ventral fronto-striatal circuitries over the course of the disease. The ventral tier which is severely degenerated in PD sends DA projections primarily to the dorsal striatum, which projects to relatively restricted portions of the more dorsal and lateral parts of the PFC (Alexander et al., 1986). The relatively intact dorsal tier sends DA projections primarily to the ventral striatum, which projects via the output nuclei of the basal ganglia and the thalamus to ventrolateral and ventromedial PFC. Reproduced from (Cools 2006).*

In PD, the external appearance of the brain is usually unremarkable (Factor & Weiner 2008; Koller & Melamed 2007a) and there is no significant atrophy of the midbrain however in some cases there is frontal atrophy (Factor & Weiner 2008). However PD is more prevalent in older individuals and age is associated with anatomical shrinkages of cortical and subcortical regions (Raz & Rodrigue 2006). The cardinal pathology in idiopathic PD consists of formation of interneuronal lewy bodies, which are eosinophilic inclusion bodies, composed mainly of misfolded α-synuclein, a 140-amino-acid protein which is a normal constituent of the presynaptic apparatus (Braak et al 2004; Koller & Melamed 2007b; Pavese & Brooks 2009). The Lewy bodies can be detected in the lower brainstem before midbrain and nigral involvement. Braak et al have proposed a six point staging system for the pathological process in PD. According to this scoring system, during the presymptomatic stages (Braak stage 1-2) the Lewy bodies and Lewy neurites are confined to medulla oblongata/pontine tegmentum and olfactory bulb/anterior olfactory nucleus, as well as affecting the medullary raphe nuclei and locus coeruleus which are rich in noradrenergic neurons. In stages 3-4 the substantia nigra and other nuclear groups of midbrain and forebrain become affected, and this is usually accompanied by the appearance of clinical symptoms. Finally in stages 5-6 Lewy bodies also appear in the neocortex (Braak et al 2004).

Advanced stages of idiopathic PD are associated with grey matter volume decrease in the BG and smaller substantia nigra volumes on volumetric Region ROI analysis on T1 images. Also with the use of probabilistic diffusion tractography it has been

demonstrated that IPD patients have lower connectivity probability between substantia nigra and putamen/thalamus (Draganski & Bhatia 2010).

The essential pathophysiological characteristic of the PD state, as a result of the neuronal damage, is increased neuronal firing activity in the output nuclei of the basal ganglia (globus pallidus pars interna (GPi) and substantia nigra pars reticulata (SNr)) leading to excessive inhibition of thalamocortical and brainstem motor systems (Obeso et al 2008), with subsequent development of Parkinsonism (Pavese & Brooks 2009). In the parkinsonian state there is decreased excitation in the D1-bearing or 'direct' pathway and increased activity in the D2 dopamine receptor expressing, or 'indirect' pathway (Obeso et al 2008).

1.6.2 PD as a model of dopamine depletion - limitations

As PD is characterised by dopaminergic cell loss it provides a model of the effects of dopamine depletion in humans. However it is not only the dopaminergic system which is affected by the pathological processes occurring in PD and due to the reciprocity of the relationship between dopamine and other neurotransmitters, dopamine depletion cannot occur in isolation. The >80% loss of dopamine in the striatum in PD is accompanied by a 50% decrease in 5HT/serotnonin levels (Wilson et al 1996). Anatomical data on the 5-HT connectivity within the basal ganglia indicates that 5-HT is in a position to modulate function by interacting with dopamine systems both at the level of substantia nigra and at the striatum. It is thus likely that abnormalities in 5-HT transmission may contribute to the neural mechanisms of PD and complications associated with long-term treatment with levodopa (Hornykiewicz 1998; Nicholson & Brotchie 2002).

In fact the pigmented neurons of the primarily adrenergic locus ceruleus undergo a similar process of neurodegeneration to that seen in the SNc (Koller & Melamed 2007a). Cholinergic cell loss in the basal nucleus of Meynert and the pedunculopontine nucleus and serotinergic cell loss from the raphe nucleus also occurs in PD (Factor & Weiner 2008) as is also described in the Braak staging (Braak et al 2004). In view of the perturbation of other neurotransmitter systems it is likely that some of the deficits observed in this patient group may be due to these other neurotransmitter abnormalities.

PD is also a disease more common in older individuals, another factor which becomes important when testing this group of patients. Concurrent with the anatomical shrinkages observed in older individuals, the efficacy of various neurotransmitter systems, including the dopaminergic, serotinergic, cholinergic and noradrenergic systems are also compromised by ageing (Eppinger et al 2011). There is age related loss of both the striatal D1 (Suhara et al 1991) and D2 (Antonini et al 1993) receptors and of the dopamine transporter DAT (Erixon-Lindroth et al 2005). The average decline ranges between 5% and 10% per decade from early to late adulthood (Eppinger et al 2011). These differences lead to difficulties when comparing subjects of different ages. However, if a smaller age range is enforced on the experimental group tested, then the generalisation of results to the whole population of PD patients becomes problematic.

In addition, Parkinson's disease is a very heterogeneous disorder with wide ranging clinical phenotypes. It is therefore difficult to make inferences in these patients across groups and specifically to compare them one with another. In my first study

(chapters 3 and 4) I employed a within subject design in order to avoid the need to match different individuals with PD with each other and in the second experiment (chapter 5) I matched PD patients with controls. Although these strategies address some of the difficulties, it still remains difficult to make inferences about the wider population of PD sufferers from small samples.

Nevertheless, despite the many difficulties involved in studying PD patients, this group of patients provides the only disease model of dopamine depletion in humans and therefore by studying PD patients in complex cognitive tasks we can gain further insight into some of the cognitive difficulties these patients face and facilitate increased understanding of the possible neurobiological mechanisms behind these deficits.

1.6.3 Mechanism of action of Levodopa and dopamine agonists

In the next section I will discuss the cognitive sequelae of the treatment of Parkinson's disease. In order to provide a background for this discussion, in the following section I provide a brief overview of the mechanism of action of Levodopa and dopamine agonists.

The mainstay of treatment for PD aims to counteract the depleted dopamine levels caused by the disease, either by replenishment or by symptomatic relief of the consequences of this depletion. There are several pharmaceutical methods which can be used to this end. The gold standard of symptomatic therapy remains oral administration of Levodopa (L-dopa) (Koller & Melamed 2007b). L-dopa is the

precursor of dopamine and as opposed to dopamine itself, can easily cross the blood brain barrier. It must be administered in a formulation along with a peripheral aromatic-amino-acid de-carboxylase inhibitor (AADI) in order to minimise its breakdown in the extracerebral tissues. AADI's do not cross the blood brain barrier and therefore do not affect conversion of dopamine in the brain (Koller & Rueda 1998; Pahwa et al 2003).

Once Levodopa has crossed the blood brain barrier it is converted to dopamine and then stored in synaptic vesicles for subsequent release (Ebadi & Pfeiffer 2005). This is the case in early PD , and in the healthy brain, so that administration of levodopa at this stage is more likely to mimic the phasic physiological role of dopamine (Factor & Weiner 2008). Later in the disease progression however, with the further loss of presynaptic dopaminergic neurons and/or with excessive doses, levodopa may also be converted to dopamine in nondopaminergic neurons leading to loss of normal physiologic control of dopamine release (Factor & Weiner 2008). The antiparkinsonian effect of levodopa is predominantly due to stimulation of the D2 receptors. D2 postsynaptic receptors and presynaptic autoreceptors have almost contrasting functions, and it is thought that the activation of postsynaptic D2 receptor is what exerts the effects of dopamine on motor behaviour (Factor & Weiner 2008). Although L-dopa is the most effective symptomatic therapy, it has drawbacks including the development of motor fluctuations and dyskinesias associated with chronic use as well as neuropsychiatric disturbances (Ebadi & Pfeiffer 2005).

Another important treatment option, utilised especially in early PD or as an adjunct to levodopa, are dopamine agonists. Dopamine agonists interact directly with the dopamine receptors and different agents have different affinities for dopamine receptor subtypes, with the most commonly used agents such as ropinirole and pramipexole having mainly D2/D3 receptor affinity (Koller & Melamed 2007b). Given their different receptor affinities, it might be expected that the different agents have different clinical effects but this has not yet been clearly established in the literature (Koller & Melamed 2007b).

We manipulated these chemical treatments in order to elucidate specific mechanisms for the actions of dopamine. In our experiments, as with the majority of drug manipulations in PD patients performing tasks seeking to elicit the effect of DRT (Dopamine replacement therapy) on performance, the 'washing out' period involves a minimum of a 12 hour withdrawal from all DRT, and omission of all long acting levodopa preparations within 24 hours. As most patients are treated with adjunct agents to the levodopa it is possible that some of the effects of medication withdrawal in the studies found in the literature, and in the first two experiments described in this thesis, are attributable to therapy with these other agents which also have a longer half life, with most dopamine agonist having a half life of between 6-20 hours (Koller & Melamed 2007b). Unfortunately due to recruitment difficulties it was not possible for us to recruit patients who were on levodopa alone a factor likely to reflect current prescribing patterns.

1.6.4 Cognitive deficits in PD and effect of LDOPA on these deficits

Although initially PD was primarily recognised as a movement disorder, recent evidence points to widespread cognitive involvement even in non-demented and non-depressed PD patients with cognitive deficits being observed even in the early stages of the disease (Lees & Smith 1983) . These cognitive deficits broadly resemble those observed in patients with frontal lobe damage and include deficits such as attentional and working memory deficits (Cools 2006; Dubois & Pillon 1997; Lees & Smith 1983; Owen et al 1992). The CamPaiGN study demonstrated that 36% of PD patients performed poorly in at least one of three cognitive tasks (Mini-Mental State Examination, a pattern recognition task, and the Tower of London task) (Foltynie et al 2004) and even early, drug naïve PD patients were found to have double the rates of mild cognitive impairment when compared to controls (Aarsland et al 2009). From the follow up of the PD patient cohort in the CamPaiGN study it also became evident that dementia in PD was related to variations in tau haplotype while the frontal-executive dysfunction observed appeared more dopaminergic in basis and had a better prognosis (Williams-Gray et al 2009).

As discussed above, the dopamine cell degeneration is not uniform across the midbrain with the SNc and dorsal striatum being particularly sensitive to neuronal degeneration (Pavese & Brooks 2009). The main projection structures of the dorsal striatum include the motor and pre-motor cortices, the supplementary motor areas and the dorsolateral PFC (Alexander et al 1986). The VTA and ventral striatum with their projections to the orbitofrontal cortex (OFC), the amygdala and the anterior

cingulate cortex (ACC) and inferotemporal cortex, are less affected early in the disease (Cools 2006).

The gold-standard of treatment of PD is replacement of depleted dopamine levels with levodopa. Younger patients however are generally commenced on other types of therapy such as dopamine agonists or MAO inhibitors in order to allow the commencement of levodopa to be postponed. The reason for this is that prolonged treatment with levodopa can lead to increased motor fluctuations and dyskinesias (Factor & Weiner 2008).

It appears the relationship between dopamine and performance in many cognitive tasks, especially working memory tasks (Stuss & Knight 2002a), follows an inverted U-shaped function whereby the optimum level of performance exists at a certain level of dopaminergic stimulation and moving off that peak, either by reducing or increasing the levels of dopamine leads to worsened task performance (Cools 2006; Williams & Goldman-Rakic 1995). It has been proposed that this optimum level differs between different cortico-striatal circuits, and therefore by increasing or decreasing the optimum level in one circuit, there might be a movement away from the optimum in different circuit. For example dopaminergic modulation with L-dopa in PD patients may move subjects to their motor optimum but away from their cognitive optimum (Gotham et al 1988; Rowe et al 2008).

It has been noted in several studies that performance on cognitive tasks while ON and OFF medication varies between different individuals with PD. There are several hypotheses explaining this observation. The first, proposed by Gotham et al (Gotham et al 1988) explains the individual variation in performance to be due to

differences in individual baseline dopamine levels causing some circuits in the striatium to have adequate dopamine levels while others become putatively 'overdosed' by the DRT (Gotham et al 1988). A different hypothesis proposed by Kulisevsky et al (Kulisevsky et al 1996), spurred by the finding that patients with different responses to dopaminergic medication (stable vs fluctuating ON and OFF periods) had different performance in cognitive tasks when ON and OFF levodopa, they proposed that the performance differences reflected differences in sensitivity to plasma levodopa concentrations due to supersensitivity of neurons to dopamine receptor stimulation (Kulisevsky et al 1996).

As mentioned above, PD results in deficits across several cognitive domains, including probabilistic learning and classification tasks (Graef et al 2010; Knowlton et al 1996), with dopamine replacement therapy (DRT)having distinct effects on these behaviours.

For example, in probabilistic reversal learning tasks, medication can impair performance in reversals, specifically those that are signalled by negative feedback (Cools et al 2006; Cools et al 2001a) but can improve task switching (Cools et al 2001a). This pattern has been attributed to an appropriate dopamine replacement of the dorsal striatal circuitry involved in task switching while 'overdosing' the relatively spared ventral striatal circuitry involved in reversal learning (Cools et al 2001a; Cools et al 2007a).

In reinforcement learning paradigms when PD patients are OFF DRT, the expression of learning from positive feedback is impaired. Conversely, when ON DRT they show impaired performance in learning from negative outcomes (Bodi et al 2009; Frank

et al 2004). This behavioural pattern has been attributed to increased levels of striatal dopamine, when patients are ON their DRT, boosting prediction error signals resulting in enhanced learning from positive outcomes. By contrast, a prevention of dips in dopaminergic activity, as observed with omission of expected outcomes, is suggested to worsen learning from negative outcomes (Frank et al 2007; Frank et al 2004; Maia & Frank 2011). In other words, we see that performance in these and many other cognitive tasks can either be ameliorated or worsened by DRT, leading to a complex picture of the function of dopamine in cognition.

Chapter 2

Methods

The following chapter reviews the methodology used in the experiments described in the next chapters.

2.1 Functional Magnetic Resonance imaging (fMRI); Physical principals

2.1.1 Magnetic resonance imaging (MRI): basic principals

Magnetic resonance imaging (MRI) is a non-invasive imaging technique which utilises the magnetic properties of atoms in order to create detailed images of internal organs.

2.1.1.1 Spin and radiofrequency pulse

Protons along with neutrons and electrons compose atoms which together comprise all matter. Neutrons and protons form the atomic nucleus and different atoms have different nuclear compositions. Under normal conditions protons spin on their axis due to thermal energy. The motion of spinning generates an electrical current on the surface of the proton which in turn creates a small magnetic source when it is placed within a magnetic field. The strength of this magnetic source is called the magnetic moment. When protons have an odd numbered atomic mass (as in the hydrogen atoms) the spin results an angular momentum (Huettel et al 2009).

A nucleus must have both a magnetic moment and an angular momentum for it to have a nuclear magnetic resonance property (NMR), and if a nucleus does not have both these properties it cannot be studied using magnetic resonance. Due to the high water content in the human body, hydrogen protons which have NMR properties, are abundant. In the absence of a magnetic field the spin axis of protons are orientated randomly and tend to cancel each other out leading to a very small

net magnetisation. To increase the net magnetisation of protons a strong magnetic field is applied which aligns the axis of the protons.

When protons are in a strong magnetic field, they make a gyroscopic motion around their aligned axis' which his known as precession. The precession frequency, known as Larmor frequency, is determined by the type of nucleus. Protons can precess in either a parallel or an antiparallel state to the magnetic field, where protons in the parallel state have lower energy then protons in the antiparallel state (figure II.I).

Copyrighted Figure removed

Figure II.I: Magnetic field causes the alignment of nuclei that have the NMR property. *(a) When no external magnetic field is present protons have their spin axis aligned randomly. (b) When a strong magnetic field is introduced protons tend to align their axis either parallel or anti-parallel to the magnetic field. More of the spins will enter the parallel state resulting in a net magnetisation parallel to that of the scanners' magnetic field. Reproduced from (Jezzard et al 2003).*

Magnetic resonance techniques measure the net magnetisation of all nuclei in a volume and the magnitude of this is determined by the difference between the number of spins in the parallel and antiparallel states, the more spins in the parallel state, the bigger the net magnetisation. One way to increase the proportion of parallel spins is to reduce the temperature, though this method is impractical for human imaging. Another way of increasing net magnetization is to increase the strength of the external magnetic field.

In order to generate an MR signal however, one must pertubate the equilibrium state of the spins and then observe how they react to this perturbation. Transitions between the low and the high energy states can be triggered by the delivery of energy to the spin system. Radiofrequency coils bombard spins in the magnetic field. The distribution of spins between the low and the high energy states is altered by the delivery of this energy, favouring transitions between the more abundant (typically the parallel low energy state) to the less abundant (typically the anti-parallel high energy state), a process known as excitation. When the amount of excitation required to create equal numbers of nuclei of each energy state is delivered, the net magnetisation is flipped form the longitudinal to the transverse axis where the measurable MR signal is greatest. When the electromagnetic waves are turned off, the excitation of the nuclei stops and due to the disruption of the thermal equilibrium due to excitation, the excess spins at the high energy level must return to the lower level to restore equilibrium. When these high energy spins fall back to the low energy state they emit photons and the energy in these photons is equivalent to the difference between the two states. The changes in

magnetisation can be detected using a radiofrequency coil and the changing current in these detector coils constitute the MR signal. From the excited state there are two forms of relaxation:

T1 relaxation

The T1 relaxation time is the time taken for the spin system to lose energy and for the net magnetisation to become re-aligned along the longitudinal axis. The T1 relaxation time influences the rate at which MR images can be collected as it renews the longditudinal magnetisation so it can be excited again.

T2 relaxation

T2 relaxation is the time taken for the decay of the transverse component of the net magnetisation to occur due to spins becoming out of phase with each other, and the T2* time is the cumulative effect of T2 time plus the effect of the magnetic field inhomogeneities and is therefore shorter then the T2 time. T2 relaxation has the effect of reducing the overall net magnetisation in the transverse plane. Depending on when the image is acquired during the relaxation process determines the intensity of the image. The BOLD contrast in fMRI relies on the T2* contrast (Huettel et al 2009).

2.1.1.2 Frequency and phase encoding

In order to be able to differentiate between different structures, the spin precession frequencies need to change over space. Frequency encoding allows the construction of a one dimensional map of proton density along the gradient due to the introduction of another magnetic gradient. By applying this extra gradient there

will be extra oscillations imposed on the precessions that are ongoing, causing different precession speeds for different protons dependant on where they are in space. The information gathered from this process can be used to map distances between structures that are being imaged.

Phase encoding refers to the application of another gradient within a slice in a sequential manner to allow for the formation of a three dimensional MR image. The phase encoding gradient is applied before the data acquisition so that the spins are precessing at different rates depending on their spatial location. So when the extra frequency encoding gradient is introduced they will already differ in their current angle of precession (their phase). In fMRI data acquisition the pulse sequences are very fast and the two gradients alternate rapidly during data acquisition. By recording the signal many times following many different combinations of gradients it is possible to effectively estimate the density and distribution of the nuclei of the object being imaged (Huettel et al 2009).

2.1.1.3 Voxels

The basic sampling unit of MRI are known as voxels, or volume elements. All MRI images of the brain are in three dimensions and voxels represent a quantity of 3D data in the images. Theoretically, the best imaging technique would have a high spatial and temporal resolution while minimising the signal to noise ratio. These factors are however at odds with each other. High spatial resolution can be achieved by reducing the voxel size due to the fact that small voxels allow for better anatomical localisation. This however comes at a price due to the fact that the total signal recovered from a voxel is proportional to its size. Therefore if voxels are too

small there may be insufficient signal to create high quality images thereby degrading the signal to noise ratio. A further complicating factor is as the voxel size gets smaller the total amount of voxels in the image increases thereby lengthening the imaging time required to collect the necessary information which in turn will affect the temporal resolution of the data collected (Buxton 2009). The voxel size acquired in the first experiments in this thesis (chapters 3 and 4) was 2x2x2mm and in the third (chapter 6) are of the order of 3x3x3mm.

2.1.1.4 Image contrast

In order to differentiate tissues from one another, it is possible to exploit the different relaxation time properties present in different tissues. Other methods for achieving this goal include injection of contrast agents, relying on changes in physiology altering magnetic properties or sensitising the images to blood flow or molecular diffusion. Here I will discuss only the contrast achieved by exploiting the differences in relaxation times. By altering two of the sequence timing parameters, the time between radiofrequency pulses (TR) and the time to echo following the excitation pulse (TE), which is the time taken for the transverse magnetisation to decay, it is possible to achieve relaxation time contrast. By shortening the TR so it is less then the T1 relaxation time (the time it takes the longitudinal magnetisation to completely recover) it is possible to achieve T1 weighted images. For example at 3 Tesla, fluid has a long T1 relaxation time (over 3sec) whereas white matter has a much shorter one (approximately 800ms) giving an image where fluid is dark and white matter a light grey. By manipulating the TE time, regions which lose their transverse magnetisation quickly will have a lower signal whereas regions with

longer T2 relaxation time such as fluid, will appear brightest. For a more extensive review see (Jezzard et al 2003).

2.1.2 Functional magnetic resonance imaging (fMRI): basic principals

2.1.2.1 BOLD contrast in fMRI

The metabolic consequences of neuronal activity provide the BOLD (Blood-oxygenation-level dependant) contrast used in fMRI. Neurons which are active require more oxygen to metabolise increased amounts of glucose, and consequentially this will cause an increase in levels of deoxygenated blood. Oxygenated and de-oxygenated blood, have different magnetic properties. When haemoglobin is oxygenated it is diamagnetic (weakly repulsed from a magnetic field), and when it is deoxygenated it is paramagnetic (attracted to a magnetic field). In the early 1980's it was experimentally verified by Thulborn et al (Thulborn et al 1982) that when the magnetic field strength was high (over 1.5 Tesla) there were differences between the transverse relaxation of oxygenated and deoxygenated blood.

This fact was exploited in seminal work by Ogawa et al (Ogawa et al 1990a; Ogawa et al 1990b), when it was demonstrated that by using gradient echo techniques it was possible to accentuate the susceptibility effects of deoxyhaemoglobin in venous blood. Since those initial experiments it has become evident that the BOLD response depends not only on blood oxygenation but on cerebral blood flow and volume. For a full review see (Logothetis 2003).

2.1.2.2 Neurophysiology of BOLD signal

The MR signal change triggered by neuronal activity is known as the haemodynamic response, and this response is dependant on the amount of deoxygenated haemoglobin over time. It is shaped not only by the extraction of oxygen by active neurons but also on changes in blood flow and blood volume. When neuronal activity increases, so does cerebral blood flow, providing an increase in the amount of oxyhaemoglobin which is actually greater than that which is required for neuronal activity. As the proportion of deoxyhaemoglobin compared to oxyhaemoglobin decreases due to the increase in cerebral blood flow, the BOLD fMRI signal increases, therefore the BOLD signal reflects the increase in cerebral blood flow to the active neuronal area rather than a direct measure of deoxygenated blood. Data from animal work in which simultaneous fMRI and neurophysiological recordings were undertaken have demonstrated that the fMRI BOLD signal correlates most closely with local field potential readings, and has been interpreted as meaning that BOLD more closely reflects inputs to an area rather than neuronal spiking activity (Logothetis et al 2001) .

The spatial resolution of fMRI studies is generally good, though it depends in part on voxel size. As mentioned earlier, the larger the voxel size the worse the spatial resolution. However if voxel size is decreased too much, the signal to noise ratio will go up, and as well as this the slice acquisition time will need to increase. However, BOLD is an indirect measure of neural activity and therefore another factor affecting spatial resolution in fMRI studies is the differences in vasculature between different areas. Because of changes in blood flow, when large amounts of

oxygenated blood flow into areas with increased neuronal activity, the proportion of oxygen extracted by neurons decreases causing the amount of oxygenated blood entering the draining venous system to increase which in turn can cause draining vessels which may be relatively distant from the area of neuronal activation, to show BOLD signal changes.

Another factor affecting spatial resolution is the variability in individual anatomy between different subjects. The preprocessing stages detailed below increase the functional resolution of the data but at the same time reduce the spatial resolution.

fMRI has an intermediate level of temporal resolution, less then that afforded by recordings from microelectrodes which can record the firing of a single neuron as it occurs but better than afforded by other approaches such as PET or lesion studies. Although cortical neuronal responses occur within tens of milliseconds following a sensory stimulus, the haemodynamic change does not commence until 1-2 seconds later. The fMRI BOLD haemodynamic response rises and falls over a period of approximately 10 seconds and so we need to estimate neuronal activity based on these slower changes in the vascular system (figure II.II).

Copyrighted Figure removed

Figure II.II: The BOLD impulse response. *The peak occurs at 4-6 seconds followed by an undershoot of 10-30 seconds. In high magnetic fields an initial undershoot may also be present. Image taken from (Henson 2008)*

2.2 Functional Magnetic Resonance imaging (fMRI); Statistical analysis

2.2.1 Pre-processing

During an fMRI experiment, data from the 3D matrix of voxels is resampled over time so that every voxel will have numerous time points. Over time however, there will be head movements which can have a large influence on the data acquired and cause difficulty when trying to align voxels with each other, which is a necessary step in order to be able to make statistical inferences about the data as a whole. It is therefore necessary to apply preprocessing algorithms to the data in order to remove uninteresting variability prior to formal statistical analysis.

Below is an image depicting the different stages of preprocessing of fMRI data (figure II.III). The stages involved in preprocessing are in the red box below. In the first stage, the images are realigned to adjust for movements between slices. Following this, in the functional-structural co-registration stage the functional and structural images are linked by overlaying one on the other. The next stage is called normalisation in which the images are warped until they are in standard MNI space (which is an average template over many normal brains) and finally the images are smoothed in order to increase signal to noise ratio. Further details of each stage are provided below.

Figure II.III: The stages of fMRI data analysis. *The stages involved in preprocessing are in the red box. The stages in which the data is modelled, details on which are below, are in the Green box. Reproduced from (Friston et al 2006)*

2.2.1.1 Spatial realignment

The aim of this stage is to correct for motion artefact by adjusting the images so that they are always in the same position. The realignment procedure consists of two steps – registration and then transformation. In the registration step, successive image volumes are realigned to a single reference volume. This is done with the use of a rigid body transformation which makes the assumption that because the size and shape of the images to be aligned are identical, they can be superimposed upon each other by using a combination of three translations (moving the whole image volume along x, y and z axis) and three rotations (rotating the image volume through the x, y and z axis) (Huettel et al 2009). In the transformation/interpolation step the data is re-sampled in order to estimate the values that would have been present had no head motion taken place. After realignment, which deals with any linear shifts, there are still significant levels of variance resulting from subject movement within the scanner therefore another step is required – unwarping. The images acquired in the scanner are distorted due to inhomogeneities in the magnetic field, these also change as the subject moves and with the use of the subject movement parameters from realignment, the field maps which assess distortion of the reference image, and estimations of changes in the magnetic field due to subject movements via iteration then an estimate of distortion at each time point can be given which can be used for unwarping of the data. Following this, the functional-structural coregistration step takes place in which the functional and structural images are aligned with each other in order to allow for the overlay of functional activations on to individuals' own anatomy.

2.2.1.2 Spatial normalisation

In this stage the images are warped so they fit the standard template brain. Different subjects will have different brain volumes and shapes and therefore in order to make statistical inferences over groups it is essential that the images can be compared with each other. Normalisation attempts to compensate for these differences in shapes by warping the images so they are the same as those of other brains. I used the unified segmentation algorithm available in SPM to perform normalisation.

2.2.1.3 Spatial smoothing

In this step the data is filtered with a Gaussian filter. Gaussian filters have the shape of a normal distribution (or bell-curve). The reason for the introduction of this filter is to spread the intensity of each voxel over nearby voxels. As fMRI data has spatial correlations due to functional similarities between adjacent brain areas and blurring due to vasculature, by smoothing data, activation is distributed over a range of voxels. Due to the fact that subjects are unlikely to have activations in exactly the same voxels as each other when the data is combined, the activation is spread over adjacent voxels and by introducing a Gaussian filter the signal to noise ratio is increased. Another advantage of smoothing is that it improves the validity of statistical techniques by increasing the normality of data because averaging over multiple observations tends towards the normal distribution and parametric statistical tests assume that error is normally distributed.

2.2.2 General linear model (GLM)

The general liner model is used to construct the design matrix which is a mathematical representation of the experiment. The core concept of this regression model is that the value of the observed data (y) can be explained by a linear combination of a set of regressors (x_i) each with a variable weighting (β_i) and residual noise in the data (error – ε). The letter β represents how much each regression factor contributes to the overall data. β_0 is the term for all the factors held constant throughout the experiment. The basic formula for the regression analysis is as follows:

$$y = \beta_0 + \beta_1 x_1 + \beta_2 x_2 + \ldots + \beta_n x_n + \varepsilon$$

In the general linear model the statistical significance of a regressor is assessed by the amount of variability it explains compared with the amount of variability explained by the error term.

2.2.2.1 Parameter estimation using GLM

The design matrix is constructed by the experimenter based on the hypothesised effects of experimental manipulations whereas the parameter weights and residual error are calculated during the analysis (figure **II.IV).**

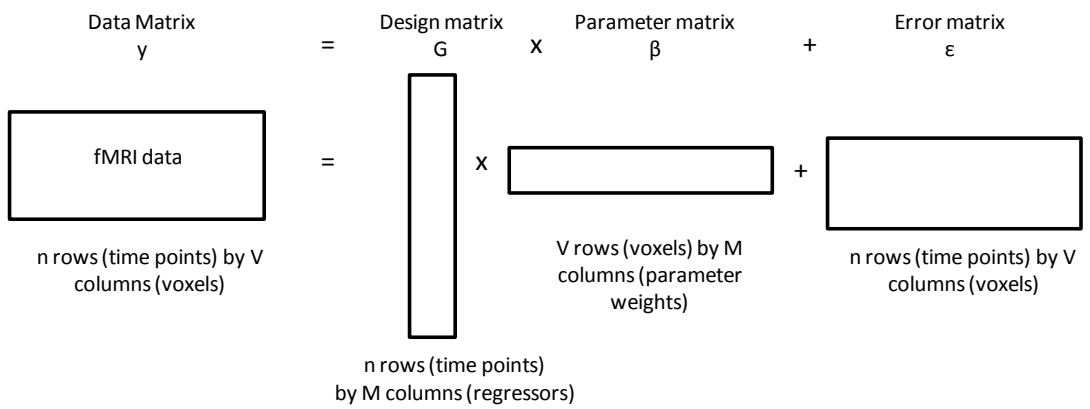

Figure II.IV: *fMRI data (y) is a two dimensional data matrix consisting of n time points by V voxels (the spatial structure is not used in the GLM they are arranged along one dimension for ease of calculation). The design matrix consisting of M regressors each n time points in length. The parameter matrix contains M parameter weights and V voxels such that each cell indicates the β value for a given voxel. The error matrix expresses the residual error for each voxel so is a n by V matrix. Reproduced from Huettel (Huettel et al 2009)*

The best fitting mode is then found classically using the sum of least squares method, although other methods are also available such as variational Bayes.

fMRI provides information about changes in activations over time and hypotheses require comparison of activations between conditions. In order to find the variability explained by a regressor of interest, contrasts are run on the matrix. In order to test an experimental hypothesis it is necessary to test whether the experimental manipulation has caused a significant shift in the parameter weights (on how much each regressor contributes to the overall data). T-contrasts assume a

directional hypothesis and are uni-dimensional (vectors) whereas F-contrasts are non-directional and multi-dimensional (matrices). There are several types of T-contrasts which can be run to establish this. The first and simplest contrast is one in which the effect of one regressor causes a difference in fMRI BOLD signal. This can identify voxels in which activation is increased compared with the other regressors (by assigning a weight of +1 to the regressor of interest and 0 to the other regressors) or decreased compared with other regressors (by assigning a weight of -1 to the regressor of interest and 0 to the other regressors). In the second type, two conditions are directly contrasted by giving one regressor the contrast weight of +1 and the other of -1. This uses subtractive logic which compares 2 variables which are assumed only to differ in one property- the independent variable. Another option is to test a non-directional hypothesis by comparing a set of contrasts with an F-test. As well as being non-directional the F-contrast can test whether any combination of contrasts can explain the variablility in the data but does not indicate which of the contrasts are driving significance, just that there is significant difference between the conditions.

2.2.3 Haemodynamic response function (HRF)

To create design matrices the haemodynamic response is convolved with predicted neural activity. There are several types of haemodynamic response functions available in different statistical packages. In the experiments described the heamodynamic response is convolved with the canonical HRF which is a mixture of gamma functions (Friston et al 2006).

It is also possible to include additional regressors in the design matrix to model additional small differences in haemodynamic onset or the shape of the haemodynamic response. This is possible with the use of temporal or dispersion derivatives which improve robustness of a model to small variations in the timing of the haemodynamic response or in the width of the haemodynamic response respectively (Friston et al 2006), these have not been used in the experiments in this thesis.

2.2.4 Event related design

There are two main experimental designs. The blocked design, in which experimental conditions are separated into distinct blocks which are presented for an extended period of time, and event-related design, consisting of short duration events often with randomised timing and order (Huettel et al 2009). The experiments in this thesis used an event related design. In event related design the fMRI data is time-locked to stimulus presentation and then averaged over all trials. The main advantage of this type of trial is that it provides better information on the shape and timing of the haemodynamic response. Other benefits include increasing the ability to randomise the trials so performance is not systematically influenced by performance on previous trials (Josephs & Henson 1999). It also reduces predictability present in blocked design allowing for the studying of aspects such as novelty and priming.

2.2.5 Inferences about subjects versus populations: Random vs Fixed effects analysis

Fixed effect analysis allows for inferences about the group of subjects who took part in the study but in order to make inferences about populations one needs to use a random effects analysis. In a fixed effects analysis, the assumption is made that the experimental effect is constant across subjects and that any inter-subject variability is due to the influence of random noise. Therefore, in this type of analysis the error term conflates within and between-subject variance. The major disadvantage with this type of analysis is that it is very sensitive to extreme results even if they are present in only a few subjects within the sample.

In order to make inferences about a population, subjects need to be treated as random variables. The most widely used method for fMRI analysis is to use mixed models whereby experimental factors are fixed but the subject factor is random. In SPM this is achieved by a two stage process. In the initial stage, contrasts of parameters are estimated from a fixed effects model for each subject and then images of these contrasts become the data for a second design matrix.

2.2.6 Problem of multiple comparisons

This problem arises from the fact that when large amounts of data are collected, as is the case in an fMRI experiment, one must perform multiple statistical tests. However the more statistical tests are performed the more likely the probability of a false positive result. There are several methods to address the problem of multiple comparisons. One of the most stringent, involves controlling for the family wise error rate (FWE) by using the Bonferroni correction. The Bonferroni correction

states that in order to correct for multiple statistical tests, the significance level of each individual test needs to be equal to the overall significance level aimed for (usually less than 0.05) divided by the total amount of statistical tests undertaken. So in fMRI data this would involve dividing the target alpha value by the total number of voxels being tested.

$$\alpha_{bon} = \alpha/V$$

where V=number of voxels and α = desired significance level for the whole family of tests.

However this correction is very stringent and although it reduces the chance of a Type I errors (false positives) it also greatly increases the chance of a Type II errors (false negatives). Furthermore, the Bonferroni correction assumes independent voxels but due to the fact that fMRI data is spatially correlated, and becomes even more spatially correlated after spatial smoothing, the Bonferroni correction is too conservative for brain images (Sarty 2007).

An alternative approach is called the false discovery rate (FDR) and this describes the probability of having at least one false positive result given the set of reported positive results. This method of correction uses the amount of suprathreshold voxels and controls the expected proportion of false positives within this group. This method is dependant on the distribution of p-values and is therefore dependant on the amount of significant activations observed in the data. This method is less stringent than the FWE correction. However, if the null hypothesis is

in fact true and there are no activations, this method will be just as stringent as other methods (Friston et al 2006).

Another way to approach the problem of multiple corrections is to do what is called a small volume correction. By defining a specific anatomical region of the brain the analysis will be restricted to many less voxels then for the whole brain analyses which will in turn lead to a much less severe correction factor. This type of correction should only be implemented based on strong a-priori hypotheses .

A further method uses thresholding based on activation clusters. Instead of assigning a p-value to each voxel, clusters of voxels are created on the basis of some initial threshold and then each cluster is assigned a p-value (Smith 2004) thus reducing the amount of statistical tests necessary and increasing the signal to noise ratio within the unit (Heller et al 2006). Part of the reasoning behind this method is that the likelihood of false positive decreases with increasing cluster size due to a reduction in the cumulative probability of two adjacent voxels being activated by chance. The main problem with this system is that it might ignore small but important activations and it also assumes that activation areas are spherical which is not always the case. Perhaps the most important difficulty with this method is that it assumes that adjacent voxels are entirely uncorrelated however due to spatial correlation in fMRI data, significant voxels tend to cluster together even if they result form noise processes (Jezzard et al 2003).

Finally a widely used method for reducing the number of independent statistical tests is to use a region of interest analysis. Regions of interest can be defined either anatomically or on the basis of previously reported lesion or functional imaging

studies. Regions may also be defined using a main effects contrast from which a region is identified and then subsequent orthogonal interaction contrasts can be tested within the region (Faro & Mohamed 2010).

When correcting for multiple comparisons an important problem arises due to the significance values often being correlated across adjacent voxels. Many noise sources such as head motion will change intensity of voxels within a brain region in a systematic fashion. Preprocessing steps correcting for motion and smoothing contribute further to this spatial correlation between activated voxels. Due to the interdependence between adjacent voxels, when corrections are made based on the total number of voxels the number of independent spatial units is greatly overestimated and the alpha value is too conservative. Therefore random field theory is required which estimates the number of statistical tests needed based on the spatial correlation/smoothness of the data. Random field theory utilises the Euler characteristic (EC), which in fMRI data is the number of blobs in an image after thresholding has takern place. In order to calculate the EC at different thresholds is is necessary to know the amount of resolution elements or resels in an image which can be calculated with the help of the full-width half-maximum Gaussian kernel (FWHM) which is applied to the imaging data during smoothing. The EC, which can also be applied to 3D fMRI data, allows the experimenter to select the thresholding level which will yield only blobs which have a less then 0.05 chance of having occurred by chance (Frackowiak et al 2003). By using this theory, the threshold for significance will be much less then the one from the Bonferroni

correction leading to fewer Type II errors (false negatives) with only minimal additional risk of additional false positives (Huettel et al 2009).

2.2.7 Computational modelling in fMRI

Over recent years, fMRI analysis has moved from simply measuring the brain responses to factors manipulated in the experimental design to experimenters trying to explain the data in terms of optimisation of brain responses by making simplifying assumptions of how the brain works with the use of mathematical models (Friston & Dolan 2010; Sommer & Wichert 2003). The central idea behind computational modelling is that there first needs to be a model which describes a mapping between a set of stimulus inputs and a set of behavioural responses. The parameters of this model, the internal operations, are then correlated with the neuroimaging data (O'Doherty et al 2007).

The approach consists of firstly fitting the computational model to subjects' actual behaviour in order to find specific values for the parameters in the model which minimize the difference between the model predictions and the behavioural data. Once the best fitting model parameters have been found, they can be regressed against the fMRI data and convolved with the haemodynamic response function. (O'Doherty et al 2007).

In the experiment described in Chapters 3 and 4 a simple prediction error based reinforcement-learning (RL) model (Sutton 1998) was used to predict a trial-by-trial measure of stimulus value (see below), and thus an outcome prediction error δ, which was defined as the difference between the actual observed outcome R (correct/incorrect = 1/0) and the current expected value of the chosen stimulus.

For each pair of stimuli A and B, the model estimated the expected values of choosing A, (Q_A) and choosing B (Q_B), on the basis of the individual sequences of choices and outcomes. The expected values were set to zero before learning. After every trial t>0 the value of the chosen stimulus (say A) was updated according to the rule $Q_A(t+1) = Q_A(t) + \alpha * \delta(t)$. The outcome prediction error is defined as the difference between the actual and expected outcome, $\delta(t)=R(t) – Q_A(t)$ with the actual outcome being either 'Correct' or 'Incorrect' (1 or 0).

Given the expected values, the probability (or likelihood) of the observed choice was estimated using the softmax rule :

$P_A(t)=\exp[Q_A(t)/\beta]/\{\exp[Q_A(t)/\beta]+\exp[Q_B(t)/\beta]\}$. The parameters α (learning rate) and β (temperature) were adjusted to maximise the likelihood of the actual choices under the model, for all subjects.

The outcome prediction errors which had been estimated by the model on a trial-by-trial basis were then used as parametric regressors in the imaging data.

2.3 Drug experiments

Two of the experiments described in this thesis involve a drug manipulation. In the experiment described in chapters 3 and 4 'Dopamine and performance in a reinforcement learning task – evidence from Parkinson's disease', a crossover design was employed and in the experiment described in chapter 6 'Expectations and violations: Probing the role of dopamine in set shifting' a double-blinded crossover design was employed.

2.3.1 Crossover design used in: "Dopamine and performance in a reinforcement learning task – evidence from Parkinson's disease"

In a three period crossover design subjects are randomised to one of six sequences (in our study OFF-OFF then OFF-ON then ON-ON medication, OFF-OFF then ON-ON then ON-OFF medication, OFF-ON then OFF-OFF then ON-ON medication, OFF-ON then ON-ON then OFF-OFF medication, ON-ON then ON-OFF then OFF-OFF medication, ON-ON then OFF-OFF then OFF-ON medication) with each subject having the same probability of being selected for each sequence.

2.3.2 Double-blinded crossover design used in: "Expectations and violations: Probing the role of dopamine in set shifting"

The medication manipulation design employed in this study was that of a randomised double-blinded crossover trial. In a two period crossover design, (such as was employed in this experiment) subjects are randomised to one of two sequences (placebo-then-drug or drug-then-placebo) with each subject having the same probability of being selected (50%) for each sequence.

2.3.3 Advantages and disadvantages of crossover designs

The main advantage of a crossover design is that each subject serves as their own control and as such sample sizes can be smaller (Gallin & Ognibene 2007). It also is a very useful method to employ when the group being studied is a heterogeneous patient group such as Parkinson's disease patients as it can significantly reduce the sample size needed to test. The main disadvantage in a crossover design is the potential for either drug or performance carryover effects. In order to try to minimise these carryover effects the sessions were scheduled a minimum of one week apart, the aim of which was firstly to attempt to minimise as much as possible any potential effect of the preceding drug manipulation on performance in subsequent sessions, by allowing for complete drug washout and also to allow the subjects to forget as much as possible their previous performance. In addition, all the symbols were changed between sessions to minimise learning carryover between the sessions. Finally, the the order of the crossover was randomised to minimise the order effect.

In the final experiment in the thesis (chapter 6) both the subject and the investigator were 'blinded' to the sequence order which has the effect of reducing investigator and subject bias (Gallin & Ognibene 2007).

Chapter 3

Dopamine and performance in a reinforcement learning task – evidence from Parkinson's disease

Behavioural results

In the experiment described in this chapter I used the model afforded by PD to tease apart the influence of dopamine on learning from its influence on action performance. I sought to dissociate dopaminergic effects on learning from effects on choice and to test whether reinforcement learning impairments in Parkinson's disease (PD) are related to the acquisition or to the performance component of reinforcement learning. In this chapter I will describe the behavioural procedure and behavioural results and in Chapter 4 I will details the fMRI procedure and the imaging results.

3.1 Introduction

Dopamine is strongly implicated in reward signalling, and plays a central role in reward learning in animals (Bayer & Glimcher 2005; Schultz 1998; Schultz et al 1997; Wise 2004; Wise & Rompre 1989) and humans (Pessiglione et al 2006). Accumulating evidence from pharmacological interventions in healthy subjects (Pessiglione et al 2006) and patients with Parkinson's disease studied ON and OFF medication (Frank et al 2007; Frank et al 2004) indicates that manipulating dopamine transmission in humans influences reward-related reinforcement learning and decision-making. A common assumption arising from these data is that dopamine exerts a direct effect on instrumental learning, a form of learning that links actions and their outcomes. Indeed, at a mechanistic level, activity in dopaminergic neurons is known to express a prediction error believed to mediate learning and updating the reward value of predictive stimuli (Schultz et al 1997). The idea that prediction error based learning is computationally implemented via activity patterns within the dopaminergic system is supported by a substantial body of experimental work across species (see (Haber & Knutson 2010) for review).

However, dopamine does not solely impact on reinforcement learning. Evidence now points to a much broader range of influences including a contribution to the control of Pavlovian approach behaviour (Dreher et al 2007; Ikemoto & Panksepp 1999; Parkinson et al 2002) as well as in motivational engagement and vigour (Bardgett et al 2009; Berridge & Robinson 1998; Boureau & Dayan 2011; Lex & Hauber 2010; McClure et al 2003b; Niv 2007; Niv et al 2007). These influences on behaviour are distinct from learning (Yin et al 2008), even in cases where they

arise from a signal that actually reports a prediction error (McClure et al 2003b). We note here that previous studies investigating the effect of dopamine typically cannot distinguish between action learning and action performance (Cools et al 2001a; Frank et al 2007; Frank et al 2004; Pessiglione et al 2006). Consequently, it remains possible that the reported dopamine-dependent impact arises not from an influence on learning as such, but rather from a modulation of the expression of learning, i.e. an effect on actual choice behaviour or performance.

Parkinson's disease (PD), is a common neurological disorder characterised by neuronal loss in the substantia nigra (SN)(Edwards 2008) that leads to depleted levels of striatal dopamine (Koller 2007). PD results in deficits across several cognitive domains, including probabilistic learning and classification tasks (Graef et al 2010; Knowlton et al 1996), with dopamine replacement therapy (DRT) having distinct effects on these behaviours. For example, when PD patients are OFF DRT, their expression of learning from positive feedback is impaired (Frank et al 2007; Frank et al 2004) and when ON DRT they show impaired performance in learning from negative outcomes (Frank et al 2004; Rutledge et al 2009). This behavioural pattern has been attributed to increased levels of striatal dopamine when patients were ON their DRT boosting prediction error signals resulting in enhanced learning from positive outcomes. By contrast, a prevention of dips in dopaminergic activity, as observed with omission of expected outcomes, has been suggested to worsen learning from negative outcomes (Frank et al 2007; Frank et al 2004; Maia & Frank 2011).

Here, I sought to dissociate dopaminergic effects on learning from effects on choice by performing neuroimaging during a reinforcement learning task with patients suffering from Parkinson's disease (PD). I employed a two stage learning task which involves separate phases of (a) acquisition and (b) a subsequent performance testing involving generalisation. This task has provided an effective means of examining the neural mechanisms underlying cognitive deficits in PD (Frank et al 2007; Frank et al 2004). These previous studies focused on learning, while here we probed the effect of dopaminergic status (ON medication, and OFF medication) both on *learning* action contingencies and on *performance* during behavioural extinction. Crucially, this dissociation between learning and performance has not been explicitly explored in previous human investigations. An influence of medication on the acquisition phase would provide support for an effect of dopamine on learning while an effect confined to performance would support the alternative hypothesis, namely that the influence of dopamine includes other processes, such as Pavlovian approach or performance motivation.

3.2 Materials and methods

The study and its procedures were approved by the National Research Ethics Service, The Joint UCL/UCLH Committees on the Ethics of Human Research (Committee A).

3.2.1 Participants

Fourteen early- to moderate-stage [H+Y stage- mean (SE) 1.69 (0.26)] idiopathic Parkinson's disease patients (ten males) (as per UK Brain bank criteria) aged between 44 and 81 years [mean (SE) 61.8 (3.3) years] participated in and completed the study. Patients were recruited from the movement disorder clinic at the National Hospital for Neurology and Neurosurgery (NHNN).

We obtained written informed consent from all subjects and transport costs were reimbursed.

Subjects were interviewed for psychiatric and neurological history as well as current and past medication. They were also examined by a clinician and asked to complete several questionnaires including a health questionnaire, a mini-mental state examination and an impulse control disorder screening questionnaire (Appendix; table 3.5.1).

One subject had difficulty understanding the task demands, and adopted an incorrect strategy for stimulus selection whereby he explicitly believed the incorrect stimulus to be correct and continued to select it despite ongoing negative feedback resulting in significantly worse than chance performance. Data from this subject are

not included in any analyses. Another subject was excluded from the imaging analysis due to an incidental finding of abnormally large ventricles, preventing successful normalisation of this dataset to a standard coordinate space. Hence, 13 subjects were analysed behaviourally and 12 subjects were analysed in the fMRI study.

Twelve of the subjects were right-handed and one was left handed. All were fluent English speakers. The duration of Parkinson's disease varied from 1 to 10 years from the time of initial diagnosis [mean (SE) 4.9 (0.96) years]. Subjects had no history of other major neurological or psychiatric disease. Patients were all on levodopa/carbidopa combinations; eight patients were also on dopamine agonists; total daily dose of levodopa/carbidopa varied from 50/12.5mg to 1000/255 mg [mean (SE) 400/100 (74.4/18.6) mg] (Appendix: table 3.5.2). We did not recruit patients on trihexyphenidyl, benzhexol or high dose tolterodine due to possible confounding effects of high dose anti-cholinergic medication, or patients on amantadine due to its effect on multiple neuromodulators.

3.2.2 Stimuli

We used a version of the generalization task introduced by Frank et al (Frank et al 2007; Frank et al 2004). Stimuli consisted of Hiragana symbols which were presented in white fonts on a black background. Each stimulus had a different probability of being correct when selected. These probabilities ranged from 80% to 20%. In the first, or acquisition, stage of the task, the symbols were paired to form 3 sets: the 80% stimulus was paired with the 20% stimulus, the 70% stimulus was paired with the 30% stimulus and the 60% stimulus with the 40% stimulus. The sets

were presented in a randomized order. In the second, or performance, phase, along with all the training pairings, the best stimulus (the one with 80% chance of being correct) and the worst stimulus (the one with only 20% chance of being correct) were presented in novel pairings with all the other stimuli (see figure III.I for task depiction).

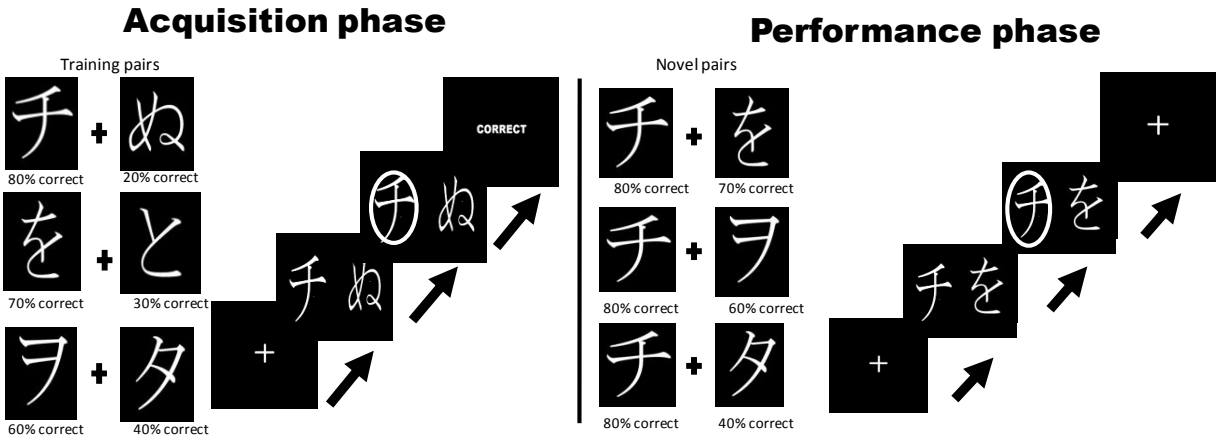

Figure III.I: Task. *Stimuli consisted of Hiragana symbols which were presented in white fonts on a black background. Each stimulus had a different probability of being correct when selected. In the first, or acquisition, stage of the task, symbols were paired to form 3 'training pairs' which remained the same throughout this phase: the 80% stimulus was paired with the 20% stimulus, the 70% stimulus was paired with the 30% stimulus and the 60% stimulus with the 40% stimulus. Subjects selected the left or right stimulus by button presses and, during the acquisition phase, also received information about the outcome (correct/incorrect). In the second, or performance, phase, along with all the training pairings, the best stimulus (the one with 80% chance of being correct) and the worst stimulus (the one with only 20% chance of being correct) were presented in novel pairings with all the other stimuli. During this phase subjects did not receive information about the outcome of their choice.*

3.2.3 Procedure

Overview

Each patient participated in three separate sessions on different days, which were a minimum of one week apart (i.e. a within subject design). Each session involved different Hiragana symbols (see figure III.I for details). All patients performed the task in three different drug states (see table 3.1 for details): acquisition and performance in the ON state (ON-ON), acquisition and performance in the OFF state (OFF-OFF), and acquisition of the stimulus contingencies in the OFF state but performance in the ON medication state (OFF-ON). The order of the different drug states in which patients performed the task was randomised. The OFF state in two of the conditions was achieved by a minimum of 12 h withdrawal from all dopaminergic medication and omission of all slow release preparations for a minimum of 18 h. On the remaining day (ON-ON), patients were asked to take their morning dopaminergic medication as usual. We were unable to test patients in the ON-OFF state, i.e. acquisition in the ON state and performance in the OFF state, due to the half life of levodopa/carbidopa combinations which would require a minimum of 7.5 hours to be metabolised and excreted resulting in too long an interval between the acquisition and performance phases. All patients were tested at similar times in the morning to equalize washout times and to control for diurnal symptom fluctuations.

	Phase 1: Acquisition	Break	Phase 2: Performance
State 1	OFF	50 mins	OFF
State 2	ON	50 mins	ON
State 3	OFF	50 mins	ON

Table 3.1: Drug states in which the task was carried out

Each patient returned three times to perform the task. In 'state 1' the subjects undertook both the first and second phases of the task in an OFF medication state. In 'state 2' subjects undertook both the first and second phases of the task in an ON medication state. In 'state 3' subjects undertook the first phase, the acquisition phase, in an OFF medication state. Following completion of the first phase, patients then received their dopaminergic medication and undertook the second phase of the task, the performance phase, in an ON medication state. The order of the states was randomised across subjects. On all three days there was a break of 50 minutes between the first and second phases of the task to allow for dopaminergic medication to be given after the first phase in 'state 3' and to allow for adequate absorption time but to ensure consistency across all 3 days.

To familiarise subjects with the structure of the task we undertook a short practice block before the first scanning session. During that practice session, patients worked on an identical task as in the main study except for the fact they were presented with different Hiragana symbols. The main session began with two functional scans (scans 1 and 2, acquisition sessions). Most subjects completed a third acquisition session on a laptop. In OFF-ON condition patients took their medication following this training. All patients then waited for 45-60mins before undergoing a third functional scanning session (scan 3, performance session) for performance testing.

On one of the three days, after the training and performance stages were complete, the patients also underwent a structural scan, a mini-mental state examination and completed questionnaires as detailed above.

Acquisition phase of the task (scanning bouts 1 and 2)

Scanning bouts 1 and 2 (acquisition phase of the task), lasted approximately 16 min, and consisted of 120 trials of 8 s each. On each trial, two Hiragana characters appeared on the screen side by side, presented via a mirror mounted on the head coil. Subjects' task was to select one of the characters on each trial by pressing either the right or the left key on a button box. The stimuli remained on the screen for 4 s, followed by presentation of the outcome (either 'Correct' or 'Incorrect') for 2 s. The likelihood of being correct or incorrect was probabilistically determined for each stimulus (see above). If subjects did not respond within the 4 s that the stimuli were on the screen the message 'no key pressed' was presented and the

trial was excluded from the analysis. A fixation cross was presented for 2 s during the inter-trial interval (ITI).

Performance phase of the task (scanning bout 3)

Scanning bout 3 (performance phase) was 10 min long and consisted of 110 trials of 6 s each. Similar to the acquisition phase, two Hiragana characters were presented side by side on each trial and subjects had to select one of the characters by pressing either the right or the left key. As before, characters remained on screen for 4 s. This time subjects did not receive feedback after making a response and the trial instead progressed immediately to the presentation of a fixation cross during the 2 s ITI.

Importantly, in addition the stimulus pairs used during training (80% with 20%, 70% with 30%, and 60% with 40%), the symbols were shown in eight novel pairings. Four of the pairings had the 'best' stimulus paired with all other stimuli (80% with 70%, 80% with 60%, 80% with 40% and 80% with 30%), and the other four pairings compared the 'worst' stimulus to all other stimuli (20% with 70%, 20% with 60%, 20% with 40% and 20% with 30%). All pairs were presented 10 times each in randomized order, resulting in 110 pairs overall (see figure III.I task depiction).

3.2.4 Data Analysis

Acquisition sessions 1, 2 and 3

All subjects reached at least 65% accuracy in the easiest pairing or after completion of 3 acquisition sessions had a minimum accuracy of 60% over all training pairs before proceeding to the performance phase. Accuracy levels in the acquisition phase were then separately computed for each drug state by averaging the overall

accuracy across all acquisition sessions on that day. The measure used was percent of trials on which the correct stimulus, i.e. the stimulus with the highest probabilistic contingency in each training pair, was selected. We then compared overall accuracy during acquisition in the ON condition to overall accuracy in the two OFF medication states using paired t-tests and a linear mixed model to detect differences in accuracy in the acquisition phase between different drug states. We also tested for differences in the acquisition rate between the different drug states by comparing learning rates in a reinforcement learning (RL) model (see below). For this test we individually fitted the parameters of the RL model to subjects' choices in the ON and OFF medication condition, comparing the resulting learning rates using a paired t-test.

Performance session

Data from the performance session were separated into trials in which the 'best' stimulus (80% chance of being correct) was presented, and trials in which the 'worst' stimulus (20% chance of being correct) was presented. We calculated the percentage of times subjects picked the best stimulus and the percentage of times the subjects avoided the worst stimulus in these novel pairings and tested for any differences in performance between the different medication conditions.

Reinforcement Learning Model

We used a simple prediction error based reinforcement-learning (RL) model (Sutton 1998) to predict a trial-by-trial measure of stimulus value, and thus an outcome prediction error δdefined as the difference between the actual observed outcome R (correct/incorrect = 1/0) and the current expected value of the chosen stimulus.

For each pair of stimuli A and B, the model estimates the expected values of choosing A, (Q_A) and choosing B (Q_B), on the basis of individual sequences of choices and outcomes. The expected values were set to zero before learning. After every trial t>0 the value of the chosen stimulus (say A) was updated according to the rule $Q_A(t+1) = Q_A(t) + \alpha * \delta(t)$. The outcome prediction error is the difference between the actual and expected outcome, $\delta(t)=R(t) – Q_A(t)$ with the actual outcome being either 'Correct' or 'Incorrect' (1 or 0). Values of stimuli that were not shown on a trial were not updated.

Given the expected values, the probability (or likelihood) of the observed choice was estimated using the softmax rule :

$P_A(t)=\exp[Q_A(t)/\beta]/\{\exp[Q_A(t)/\beta]+\exp[Q_B(t)/\beta]\}$. The parameters α (learning rate) and β (temperature) were adjusted to maximise the likelihood of the actual choices under the model, for all subjects. Trial-by-trial outcome prediction errors estimated by the model were then used as parametric regressors in the imaging data.

3.3 Results

We employed a within subject design given the inherent difficulty in accurately matching PD patients with different disease severity, we also believe that this design allowed us to minimise as far as is possible individual cognitive and genetic differences that may exist in our cohort, allowing us to look at the within subject effects of drug on behaviour. In parallel with our behavioural analysis, we also acquired neural data using fMRI. The fMRI results will be discussed in the next chapter (chapter 4). Thus, our design enabled us to explore the effect of dopamine on behaviour and on the brain by testing patients in three different drug states; acquisition and performance ON medication, acquisition and performance OFF medication and acquisition in an OFF medication state and performance in an ON medication state. This inclusion of the latter condition specifically enabled us to probe whether dopamine affected the acquisition or performance aspects of the task.

Acquisition phase results

At the end of the acquisition phase, average choice accuracy on the training pairs did not differ between groups in different drug states (paired t-tests comparing ON-ON with OFF-ON; $T_{(1,12)}$ =0.15, p=0.87, comparing ON-ON with OFF-OFF; $T_{(1,12)}$ =0.095 p=0.92, comparing OFF-ON; with OFF-OFF $T_{(1,12)}$= -.079p=0.93) (table 3.2).

Subject	ON_ON	OFF_ON	OFF_OFF
1	77%	85%	71%
2	71%	87%	87%
3	70%	71%	65%
4	70%	62%	82%
5	97%	100%	92%
6	80%	79%	86%
7	63%	53%	84%
8	56%	62%	68%
9	94%	53%	65%
10	67%	62%	64%
11	34%	70%	65%
12	92%	63%	59%
13	75%	89%	54%

Table 3.2: Choice accuracy

Average choice accuracy on the training pairs after the final training session in each of the medication groups. All subjects reached at least 65% accuracy in the easiest pairing or after completion of 3 acquisition sessions had a minimum accuracy of 60%.

Similarly, we found no significant difference in learning rates between patients when they were in an OFF compared to ON medication state (mean (SE) ON 0.25 (0.02) and OFF 0.24 (0.01), paired t-test $T_{(1,12)}$ =0.117, p=0.90) (figure III.II, table 3.3), or in the number of sessions required to reach criteria (mean (SE) ON 1.23 (0.12) OFF 1.38 (0.16) paired t-test $T_{(1,12)}$ =-0.69, p=0.50). There were also no differences between positive and negative learning rates.

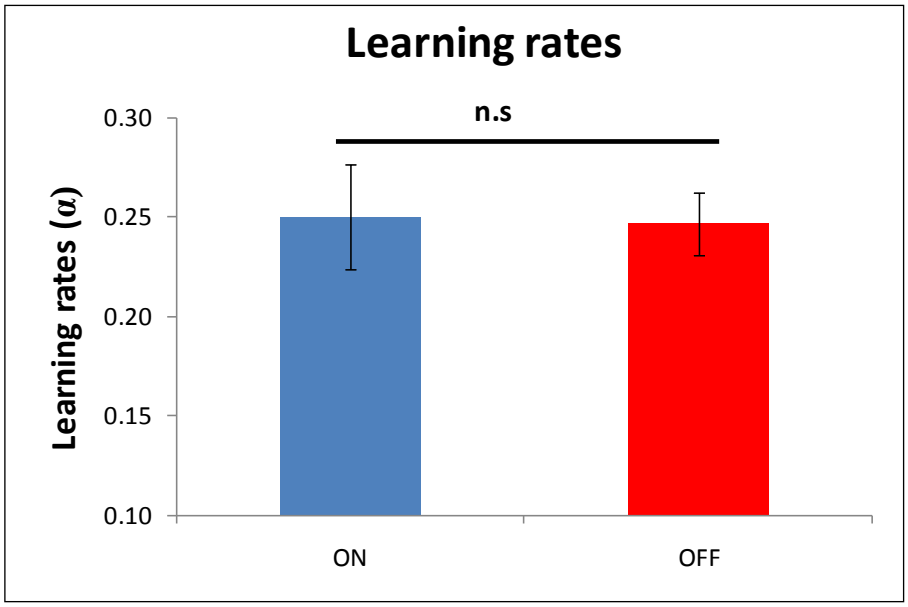

Subject	ON	OFF
1	0.175	0.3221
2	0.2324	0.2351
3	0.2267	0.2012
4	0.2797	0.3064
5	0.1486	0.1779
6	0.3678	0.1712
7	0.3443	0.2851
8	0.2004	0.2348
9	0.1935	0.219
10	0.473	0.3216
11	0.2527	0.2078
12	0.1828	0.2062
13	0.175	0.3221

Figure III.II and table 3.3: Learning rates

Learning rates were fitted separately per subject in the different drug conditions. All bouts were considered equally during fitting. The data in the 'ON' group came from the 3 bouts they performed in the ON medication state; the data in the 'OFF' group came from the 6 bouts they performed in the OFF medication state.

Performance phase results:

In the performance phase, along with all the training pairings, we presented the best (the one with 80% chance of being correct) and worst stimulus (the one with only 20% chance of being correct) in novel pairings with all the other stimuli (see above for task depiction). We found that patients ON their DRT performed significantly better than patients OFF their DRT (main effect comparing accuracy of the mean of ON-ON/OFF-ON with OFF-OFF, paired t-test, $T_{(1,12)}$ =2.8, p=0.01). Crucially, a separate examination of the three drug states revealed a main effect of drug on performance but not on acquisition (figure III.III). Subjects who acquired the contingencies in an OFF medication state and received their DRT after the acquisition phase, but before the performance phase, had the same level of overall accuracy as subjects who both acquired the contingencies ON medication and performed ON medication (paired t-test comparing ON-ON with OFF-ON, T=-0.03, p=0.97). Both the ON-ON and OFF-ON groups were significantly more accurate then the OFF-OFF group (paired t-test comparing ON-ON with OFF-OFF, T=2.17 p=0.05; and comparing OFF-ON with OFF-OFF, $T_{(1,12)}$ =2.28 p=0.04). A mixed effects linear model showed a significant effect of drug state on the performance phase ($F_{(1,36)}$=5.38 p=0.02) but not on the acquisition phase ($F_{(1,36)}$=0.002 p= 0.96).

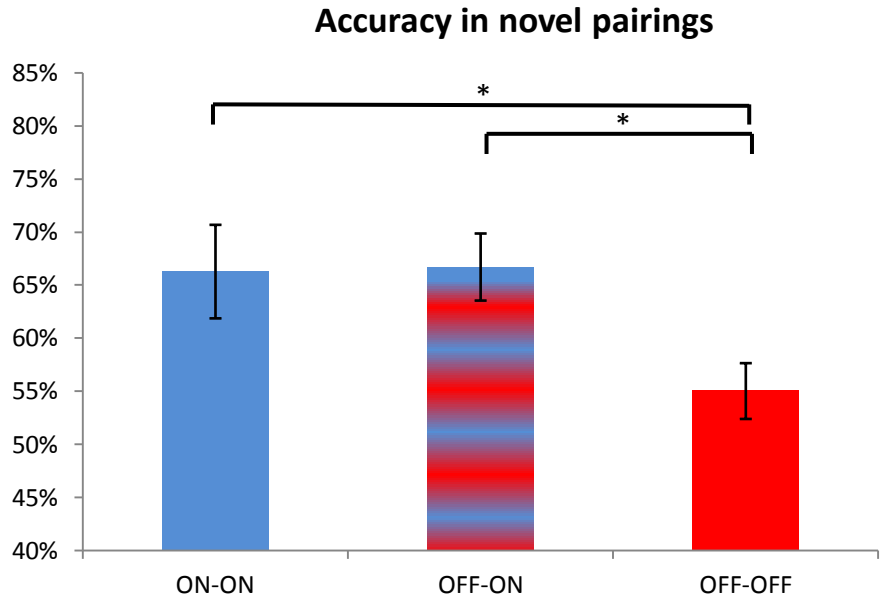

Figure III.III: Accuracy during performance phase in novel pairings. *Accuracy was significantly higher when subjects were ON DRT during the performance phase than when they were OFF (p=0.01). This effect is independent of drug state during the previous acquisition phase. Shown is the combined accuracy in selecting the best stimulus and avoiding the worst stimulus over the 3 drug states when subjects had to pick the stimulus with the highest likelihood of being correct when presented in novel parings. ON-ON session (blue) –when patients took their usual DRT, OFF-ON session (blue/red striped) when patients took their DRT only after completing the acquisition phase, OFF-OFF session (red) when patients abstained from their DRT throughout the task. The error bars represent s.e.m.*

In addition to the novel pairings, subjects were also presented with the three stimulus pairs on which they had been trained during acquisition. Interestingly, we found no differences in accuracy levels on these training pairs across the different drug states (paired t-tests comparing ON-ON with OFF-ON $T_{(1,12)}$ =-1.36, p=0.19, comparing ON-ON with OFF-OFF $T_{(1,12)}$ = -0.64 p=0.52, comparing OFF-ON with OFF-OFF $T_{(1,12)}$ = 1.26 p=0.23) (figure III.IV). There were no differences in the training pair compared with novel pair accuracy in ON-ON or OFF-ON drug states (paired t-tests comparing training pair accuracy with novel pair accuracy in ON-ON $T_{(1,12)}$ = -.93, p=0.36; and OFF-ON drug states $T_{(1,12)}$ =-.99, p=0.33), but as expected, there was a significant difference between training pair compared with novel pair accuracy in the OFF-OFF drug state ($T_{(1,12)}$ =-3.65,p= 0.003).

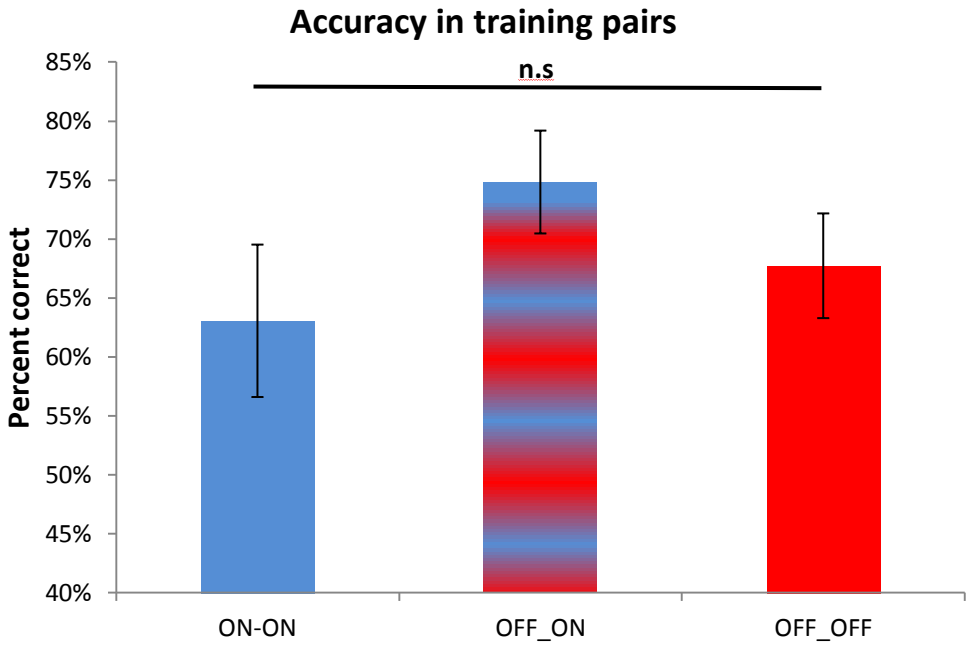

Figure III.IV: Accuracy during performance phase in training pairs. *Accuracy in selecting the better stimulus among the training pairs during the performance phase did not differ between drug states. Shown is the combined accuracy in selecting the best stimulus and avoiding the worst stimulus over the 3 drug states when subjects had to pick the stimulus with the highest likelihood of being correct when presented in the contingencies which they have previously trained on. ON-ON session (blue) –when patients took their usual DRT, OFF-ON session (blue/red striped) when patients took their DRT only after completing the acquisition phase, OFF-OFF session (red) when patients abstained from their DRT throughout the task.*

The maintained performance in the training pairs also makes an extinction effect, whereby patients OFF medication are less sensitive to the lack of positive feedback and therefore perform progressively worse during the session, much less likely. This is due to the fact that patients OFF medication were able to maintain their performance on the training pairs, which were randomly interspersed with the novel pairs, throughout the test session. This observation also makes it unlikely that the performance differences between the groups are due to a selective sensitivity in the OFF group to the time delay between acquisition and transfer.

We next tested for differential performance in selecting the best, and avoiding the worst, stimulus within the novel pairings. Interestingly, being in the ON DRT state during the performance phase selectively improved accuracy in selecting the best stimulus compared to avoiding the worst stimulus for novel stimuli pairs (paired t-tests comparing ON accuracy for picking the best compared with avoiding worse stimulus $T_{(1,12)}$ =2.16, p= 0.05). This performance difference between selecting the best and avoiding the worst stimulus was not evident when subjects both acquired and performed the task in the OFF medication state (paired t-tests comparing OFF accuracy for picked best compared with avoiding worse stimulus $T_{(1,12)}$ =0.58, p=0.56), although their overall performance was worse (figure III.V). Of note, there was no interaction between the medication status (ON vs OFF) during performance and the ability to pick the best compared with avoiding the worst stimulus as has previously been reported (Frank et al 2007; Frank et al 2004).We only found this selective improvement in picking the best stimulus compared with avoiding the worst stimulus within the ON group.

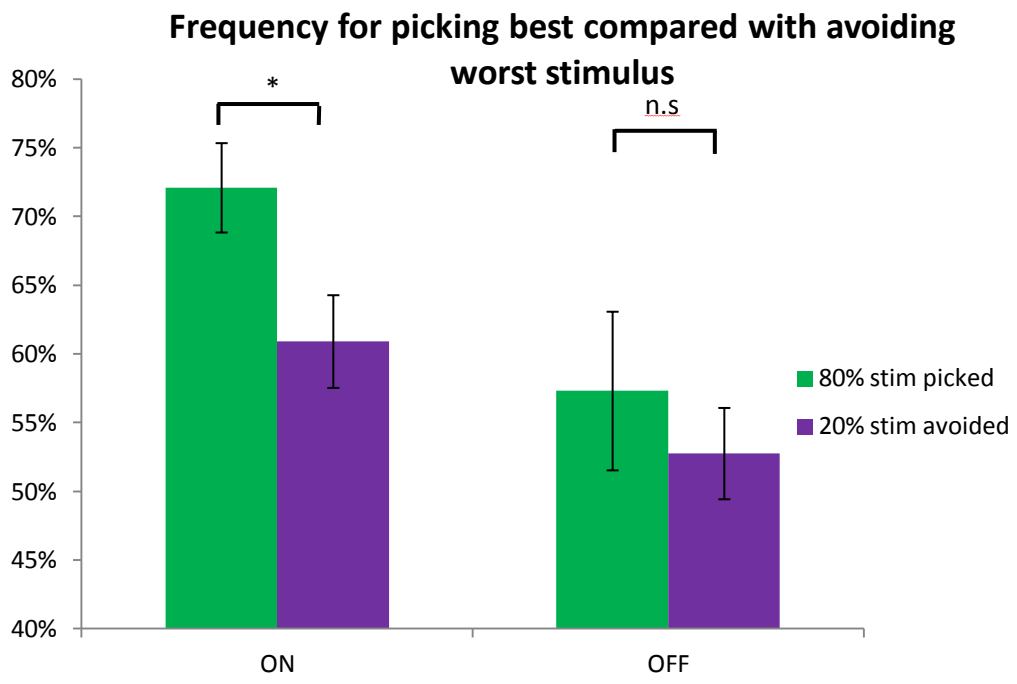

Figure III.V: Differences in accuracy at picking best compared with avoiding worst stimuli. *ON state during performance phase selectively improved accuracy for picking the best stimulus (the 80%) compared to avoiding the worst stimulus (the 20% stimulus) in novel pairings. The data in the 'ON' state comes from the 2 performance bouts performed in this medication state and the data in the 'OFF' state comes from the 1 performance bout performed in this medication state.*

3.4 Discussion

We show a striking effect of dopamine replacement therapy (DRT) on the ability of PD patients to select the highest valued stimulus in a probabilistic reinforcement learning task. Importantly, our data show that medication status at the *acquisition* task phase did not impact on successful task learning. Instead, the data show that the critical factor was medication status at the performance phase, by which time stimulus values had already been acquired successfully. The findings challenge the proposal that the impact of dopaminergic status on this form of decision making solely reflects its involvement in learning.

Our key observation was that patients who were OFF dopamine during the second task phase performed significantly worse when stimuli occurred in novel pairings. However, dopaminergic drug state did not impact on their ability to select the best stimulus when they were required to pick between pairs on which they had been trained in the first phase of the task. This indicates that the subjects OFF medication could successfully retrieve learnt contingencies but were unable to use this knowledge to make correct choices when they had to select between novel stimulus pairings. There was no difference in learning rates or accuracy during the acquisition phase between the different drug conditions, indicating that dopamine did not impact on the ability to learn stimulus values. Consequently, it would appear that DRT impacted upon the ability to generalize, in a context where subjects needed to select the best stimulus in a state characterized by novel pairings.

The fact that patients in all three drug states performed equally well when they were selecting the best cue for sets on which they had been previously trained further indicates that dopamine did not influence patients' accuracy by a direct influence on learning. Levodopa medication in PD patients has previously been shown to have a positive effect on generalisation of learnt information in novel contexts, however those observations were on a background of impaired learning and therefore crucially different from our current findings (Myers et al 2003; Shohamy et al 2006). Of course, many different systems are likely to be involved in learning, only some of which depend directly on dopamine (Beninger 1983; Daw et al 2005; Dickinson et al 2000; Palmiter 2008), and we cannot discount the possibility that a more complex learning task, such as one involving sequences of choices, might be necessary to fully reveal effects of dopamine on learning.

Conversely, beyond its putative role in learning, dopamine is implicated in a number of distinct processes related to motivation, including the control of preparatory Pavlovian conditioned responses, and motivational vigour (Bardgett et al 2009; Berridge 2007; Boureau & Dayan 2011; Dickinson et al 2000; Mazzoni et al 2007; Niv 2007; Parkinson et al 2002; Salamone et al 2003). Importantly, these remain consistent with the fact that the phasic activity of dopamine neurons codes for an appetitive prediction error (McClure et al 2003b). However, our study has enabled us to disentangle these effects from a mere effect on learning in a manner that provides clear evidence that dopamine has a specific role in action performance distinct from learning. However within this task design, we are unable to surmise whether the effect of medication was due to the boosting of tonic or

phasic dopamine levels as the drug manipulation (which included the withdrawal and then reinstatement of both L-DOPA and dopamine agonists) is likely to have had an effect on both.

A significant finding from this study is that when patients were ON their DRT, they were worse at avoiding stimuli with the poorest probabilistic contingencies than at choosing the stimuli with the best probabilistic outcomes. This is in keeping with previous research showing a similar outcome valence performance asymmetry, whereby patients ON their DRT are impaired at avoiding the least rewarding stimuli (Frank et al 2007; Frank et al 2004). It has been postulated that this worsening in performance is due to 'overdosing' of the striatum which interferes with the dips in dopamine that express negative prediction errors (Frank et al 2007; Frank et al 2004). However, in our study as in several others (de Wit et al 2011; Jocham et al 2011), we did not find a direct effect of medication on learning and we postulate that the worsened performance may reflect some other mechanism, perhaps an impaired expression of avoidance behaviour in a high dopamine state. Of note, we did not find an interaction between medication state and picking the best compared with avoiding the worst stimulus which has been reported in some previous studies (Frank et al 2007; Frank et al 2004; Voon et al 2010). We found an overall improvement in performance when subjects were ON and an asymmetry in this performance accuracy between pick best compared with avoid worst trials within this group.

By teasing apart learning and performance during a reinforcement learning task in Parkinson's patients I found that dopaminergic medication impacted performance,

but not learning. Thus, the improved performance in patients ON medication cannot be attributed to an effect on learning and must reflect some other effect of dopamine, perhaps Pavlovian appetitive approach or motivational vigour.

3.5 Appendix

Table 3.5.1: Neuropsychological data sets

	Patients (n=13)
Age	61.8 (3.3)
Education (years since age 16)	4.3 (1)
MMSE	29 (0.32)
ICD	1.6 (0.8)

Values represent mean (SE). BDI = Beck Depression Inventory; MMSE=Mini Mental State Examination; ICD = impulse control disorder questionnaire.

Table 3.5.2: Medications

	Patients (n=13)
levodopa/carbidopa	13
Stalevo (levodopa/carbidopa/entacapone)	1
Ropinirole	5
Pramipexole	2
Selegiline	2
Rasagiline	2
Anti-hypertensives	3
Anti-depressants (SSRI/SNRI)	1
Gliclazide	1
Omeprazole	1
Ceterizine	1
Detrusitol	1
Voltarol	1
Sildefinil	1
Aspirin	1

Chapter 4

Dopamine and performance in a reinforcement learning task – evidence from Parkinson's disease

Functional MRI results

This chapter follows directly from findings reported in Chapter 3 in which I described the behavioural findings that aimed to tease apart the influence of dopamine on learning from its influence on action performance. In this chapter I will describe the fMRI procedure and results. The aim of fMRI scanning was to elucidate the neural mechanisms underlying the observed differences in behaviour allowing a better understanding of the way in which dopamine exerts its influence on accuracy in a reinforcement learning task.

4.1 Introduction

As mentioned in the previous chapter, the roles dopamine plays in decision making have important theoretical, empirical and clinical implications. Here, I examined this issue by exploiting the dopamine lesion deficit model afforded by Parkinson's disease. We studied patients in a two stage reinforcement learning task, while they were ON and OFF dopamine replacement medication. In this chapter I will discuss the fMRI procedure and findings.

At a mechanistic level, activity in dopaminergic neurons is known to express a prediction error believed to mediate learning and updating the reward value of predictive stimuli (Schultz et al 1997), And the idea that prediction error based learning is computationally implemented via activity patterns within the dopaminergic system is supported by a substantial body of experimental work across species (see (Haber & Knutson 2010) for review). My first aim was to test whether dopaminergic drug state in PD patients (ON or OFF medication) had an influence on the neural representation of the magnitude of the prediction error signal during learning.

Given that dopamine does not solely impact on reinforcement learning and has a role in the control of Pavlovian approach behaviour (Dreher et al 2007; Ikemoto & Panksepp 1999; Parkinson et al 2002) as well as in motivational engagement and vigour (Bardgett et al 2009; Berridge & Robinson 1998; Boureau & Dayan 2011; Lex & Hauber 2010; McClure et al 2003b; Niv 2007; Niv et al 2007), I aimed to identify a neural basis for the behaviour we observe with the aim of differentiating the effect

of dopamine on learning, potentially mediated via its effect on the prediction error signal, from its effects elsewhere.

In the previous chapter I detailed behavioural findings which demonstrate that medication status during the acquisition phase did not impact on successful learning and instead, the data showed that the critical factor was medication status at the performance phase, by which time stimulus values had already been acquired successfully. These findings challenge the proposal that the impact of dopaminergic status on this form of decision making solely reflects its involvement in learning.

My first aim was to establish if, despite the lack of behavioural differences, whether we could identify any neural differences between the patients in different drug states during the acquisition phase of the task. My second aim was to identify the neural substrate underlying the improved behaviour in the patients who were ON medication during the performance phase of the task in the hope that we would find a neurobiological correlate for this improved behaviour.

4.2 Materials and methods

4.2.1 Experimental Paradigm

4.2.1.1 Recap of design (for further details see chapter 3)

We used a within subject design that enabled us to study the same group of PD patients (early- to moderate-stage [H+Y stage- mean (SE) 1.69 (0.26)]) in a generalization task introduced by Frank et al (Frank et al 2007; Frank et al 2004) in three separate drug states (see figure IV.I for details of task design and table 3.1 for experimental structure). We employed a within subject design given the inherent difficulty in accurately matching PD patients with different disease severity. In parallel with our behavioural analysis, we also acquired neural data using fMRI. Thus, our design enabled us to explore the effect of dopamine on behaviour and on the brain by testing patients in three different drug states; acquisition and performance ON medication, acquisition and performance OFF medication and acquisition in an OFF medication state and performance in an ON medication state. This inclusion of the latter condition specifically enabled us to probe whether dopamine affected the acquisition or performance aspects of the task (see chapter 3 for more detailed methods).

In the task, stimuli consisted of Hiragana symbols which were presented in white fonts on a black background. Each stimulus had a different probability of being correct when selected. In the first, or acquisition, stage of the task, symbols were paired to form 3 'training pairs' which remained the same throughout this phase: the 80% stimulus was paired with the 20% stimulus, the 70% stimulus was paired with the 30% stimulus and the 60% stimulus with the 40% stimulus. Subjects

selected the left or right stimulus by button presses and, during the acquisition phase, also received information about the outcome (correct/incorrect). In the second, or *performance*, phase, along with all the training pairings, the best stimulus (the one with 80% chance of being correct) and the worst stimulus (the one with only 20% chance of being correct) were presented in novel pairings with all the other stimuli. During this phase subjects did not receive information about the outcome of their choice (figure IV.I-depiction of task).

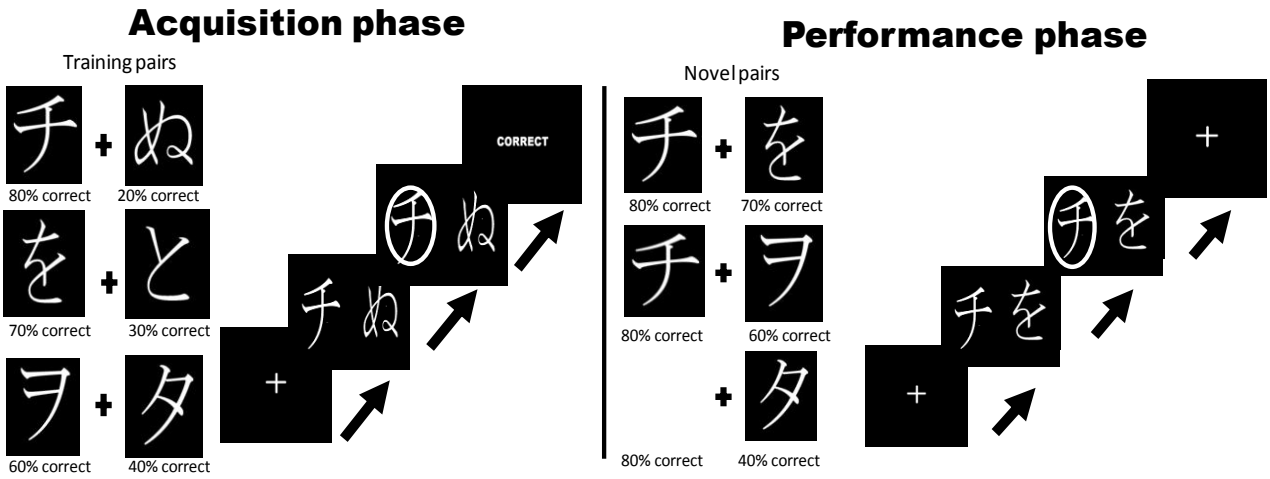

Figure IV.I: Task. *Stimuli consisted of Hiragana symbols which were presented in white fonts on a black background. Each stimulus had a different probability of being correct when selected. In the first, or acquisition, stage of the task, symbols were paired to form 3 'training pairs' which remained the same throughout this phase: the 80% stimulus was paired with the 20% stimulus, the 70% stimulus was paired with the 30% stimulus and the 60% stimulus with the 40% stimulus. Subjects selected the left or right stimulus by button presses and, during the acquisition phase, also received information about the outcome (correct/incorrect). In the second, or performance, phase, along with all the training pairings, the best stimulus (the one with 80% chance of being correct) and the worst stimulus (the one with only 20% chance of being correct) were presented in novel pairings with all the other stimuli. During this phase subjects did not receive information about the outcome of their choice.*

4.2.1.2 MRI scanning

The study was conducted at the Wellcome Trust Center for Neuroimaging at University College London using a 3T (Siemens TRIO) scanner equipped with a Siemens 12 channel phased array head coil. Anatomical images were acquired using modified equilibrium fourier transform T1 gradient echo scans, which were followed by 1-mm-thick axial slices parallel to the anterior commissure–posterior commissure plane. Functional scans used a gradient echo sequence; repetition time, 2.04 s; echo time 30 ms; flip angle 90 degree; matrix size 64 x 64; field of view 192 mm; slice thickness, 2 mm and interslice distance factor was 1mm. A total of 30 axial slices were sampled. The in-plane resolution was 2 x 2 mm.

Functional imaging data were analyzed using statistical parametric mapping software (SPM5; Wellcome Trust Centre for Neuroimaging, London, UK; http://www.fil.ion.ucl.ac.uk/spm). During preprocessing, images were realigned with the first volume (after discarding six volumes to allow for T1 equilibration effects), and unwarped. For each subject, the mean functional image was coregistered to a high-resolution T1 structural image. This image was then spatially normalized to standard Montreal Neurological Institute (MNI) space using the "unified segmentation" algorithm available within SPM5 (Ashburner & Friston 2005) with the resulting deformation field applied to the functional imaging data. These data were then spatially smoothed using an isotropic 6-mm full-width half-maximum Gaussian kernel.

Acquisition phase of the task (scanning bouts 1 and 2 and laptop session 3)

Scanning bouts 1 and 2, (acquisition phase of the task) lasted approximately 16 min, and consisted of 120 trials of 8s each. On each trial, two Hiragana characters appeared on the screen side by side, subjects completed the first two sessions in the fMRI scanner and the final session outside the scanner on a laptop in a well lit room. Subjects' task was to select one of the characters on each trial by pressing either the right or the left key on a button box while in the scanner or the left or right keyboard shift key when using a laptop. The stimuli remained on the screen for 4 s, followed by presentation of the outcome (either 'Correct' or 'Incorrect') for 2 s. The likelihood of being correct or incorrect was probabilistically determined for each stimulus (see above). If subjects did not respond within the 4 s that the stimuli were on the screen the message 'no key pressed' was presented and the trial was excluded from the analysis. A fixation cross was presented for 2 s during the inter-trial interval (ITI).

Performance phase of the task (scanning bout 3)

Scanning bout 3 (performance phase) was 10 min long and consisted of 110 trials of 6 s each. Similar to the acquisition phase, two Hiragana characters were presented side by side on each trial and subjects had to select one of the characters by pressing either the right or the left key. As before, characters remained on screen for 4 s. This time subjects did not receive feedback after making a response and the trial instead progressed immediately to the presentation of a fixation cross during the 2 s ITI.

Importantly, in addition the stimulus pairs used during training (80% with 20%, 70% with 30%, and 60% with 40%), the symbols were shown in eight novel pairings. Four of the pairings had the 'best' stimulus paired with all other stimuli (80% with 70%, 80% with 60%, 80% with 40% and 80% with 30%), and the other four pairings compared the 'worst' stimulus to all other stimuli (20% with 70%, 20% with 60%, 20% with 40% and 20% with 30%). All pairs were presented 10 times each in randomized order, resulting in 110 pairs overall.

4.2.2 Data Analysis

Reinforcement Learning Model

We used a simple prediction error based reinforcement-learning (RL) model (Sutton 1998) to predict a trial-by-trial measure of the stimulus values, and thus an outcome prediction error δ, which is defined as the difference between the actual outcome R (correct/incorrect = 1/0) and the expected value of the chosen stimulus (please see chapter 3 for full details).

Trial-by-trial outcome prediction errors estimated by the model were then used as parametric regressors in the imaging data.

fMRI analysis: whole-brain general linear model parametric analysis

Acquisition session

Functional MRI (fMRI) time series were regressed onto a composite general linear model (GLM) containing four regressors: trial onset time (the appearance of the hiragana characters), outcome onset time, motor response time and fixation cross presentation time. The outcome onset was parametrically modulated by the prediction error as estimated by the RL model (see above). We also composed

another GLM in which there were four regressors: correct trial onset time, incorrect trial onset time, motor response time and fixation cross presentation time. The actual value of the chosen cue in each trial was entered as a parametric modulator of the two trial onset regressors.

Performance session:

Four regressors were entered into the fMRI model: correct trial onset time, incorrect trial onset time, motor response time and fixation cross presentation time. The actual value of the chosen cue in each trial was entered as a parametric modulator of the two trial onset regressors.

The regressors were convolved with the canonical HRF, and low frequency drifts were excluded with a high-pass filter (128-s cutoff). Short-term temporal autocorrelations were modeled using an AR(1) process. Motion correction regressors estimated from the realignment procedure were entered as covariates of no interest. Statistical significance was assessed using linear compounds of the regressors in the GLM, generating statistical parametric maps (SPM) of t values across the brain for each subject and contrast of interest. These contrast images were then entered into a second-level random-effects analysis using a one-sample t test against zero.

Anatomical localization was carried out by overlaying the t-maps on a normalized structural image averaged across subjects, and with reference to an anatomical atlas (Naidich 2009). All coordinates are reported in MNI space (Mazziotta et al 1995).

Region of interest analysis

We extracted data for all region of interest analyses using a cross-validation leave-one-out procedure: we reestimated our main second-level analysis 12 times, always leaving out one subject. Starting at the peak voxel for the chosen cue value signal in ventromedial prefrontal cortex and Nucleus accumbens , which was identified by looking over all correct trials (in both the ON and OFF drug states), we selected the nearest maximum in these cross-validation second-level analyses. Using that new peak voxel we then extracted the data from the left-out subject and calculated a representative timecourse for each ROI as first eigenvariate from data in all voxels within a 4mm sphere around that peak.

4.3 Results

Acquisition phase results

As detailed previously, at the end of the acquisition phase, average choice accuracy and learning rates on the training pairs did not differ between groups in different drug states.

In the neuroimaging data, we examined brain responses that correlated with outcome prediction errors computed from a reinforcement learning (RL) model, fit to subjects' behavior during the acquisition phase. We found that bilateral responses in the striatum (central coordinates right putamen x=26, y=0, z=-4 left putamen x=-28, y=-12, z=-2) (figure IV.II A, B) strongly correlated with reward prediction errors, consistent with many previous results (McClure et al 2003a; O'Doherty et al 2003; Schonberg et al 2010; Schultz 1998; 2010; Schultz et al 1997). However, akin to our behavioral findings, we found no differences in prediction error related brain activation between the different drug states during acquisition (paired t-test ON compared with OFF; $T_{(1,11)}$ =-.076, p=0.46). When we examined positive and negative prediction errors separately we also did not find any differences between the different drug states (paired t-tests comparing positive prediction errors ON compared with OFF; $T_{(1,11)}$ =-.083, p=0.42; and comparing negative prediction errors ON with OFF, $T_{(1,11)}$ =-.051, p=0.614). Perhaps most surprisingly, at the time of cue onset we did not observe any correlation between brain activity and the value of the chosen cue in any of the drug states.

A

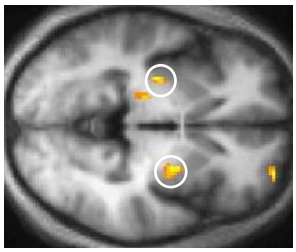

B

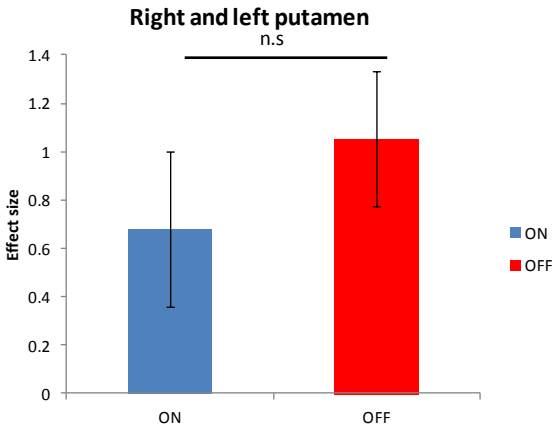

***Figure IV.II**: **Prediction error related activity during acquisition.** (A) Brain activity in putamen correlated with magnitude of outcome prediction errors across all trials during the acquisition phase. Activations are thresholded at $p < 0.001$ uncorrected. (B) Correlation (normalized effect size) between outcome prediction errors and BOLD activity in the two different drug states. Data in the 'ON' state was averaged across the 2 acquisition bouts performed in the scanner under this medication state and data in the 'OFF' state across the 4 acquisition bouts performed in this medication state. Error bars represent s.e.m.*

Performance phase results

In the performance phase we found that patients ON DRT performed significantly better than when they were OFF DRT in the novel pairings and crucially, a separate examination of the three drug states revealed a main effect of drug on performance but not on acquisition (see chapter 3).

In order to investigate for neural mechanisms underlying the observed behavioral effects during the performance phase we tested for differences in the degree at which fMRI BOLD activity correlated with decision variables between different drug states. We tested if neural representations of stimulus values at the time of cue presentation differed between drug states. We found that BOLD activity in nucleus accumbens (central coordinates x = 8, y = 12, z = -4) correlated with the value of the chosen cue, but this effect was only evident in the ON medication state for correct trials (one sample t-test, $T_{(1,11)}$ =2.7, p=0.01). Cue evoked BOLD activity did not correlate with the value of the chosen cue when patients were OFF their DRT (one sample t-test, $T_{(1,11)}$ =0.98, p=0.34) or made an incorrect choice (one sample t-test ON incorrect, $T_{(1,11)}$ =-2.12, p=0.06, OFF incorrect $T_{(1,11)}$ =-0.06, p=0.94) (figure IV.IIIA and IV.IIIB). We found an identical effect in ventromedial prefrontal cortex (x = -2, y = 38, z = 0), where BOLD activity varied with the value of the chosen cue when patients were both ON medication and made the correct choice (one sample t-test, $T_{(1,11)}$ =2.52, p=0.02), but not when they were OFF their DRT (one sample t-test, $T_{(1,11)}$ =0.31, p=0.76), or made an incorrect choice (one sample t-test ON incorrect, $T_{(1,11)}$ =-1.28, p=0.22, OFF incorrect $T_{(1,11)}$ =0.76, p=0.46) (figure IV.IIIC and IV.IIID). These findings show that activity in NAc and vmPFC successfully reflect the

values of the most rewarding cue only in an ON medication state, a characteristic that precisely mirrors patients' improved performance in this state.

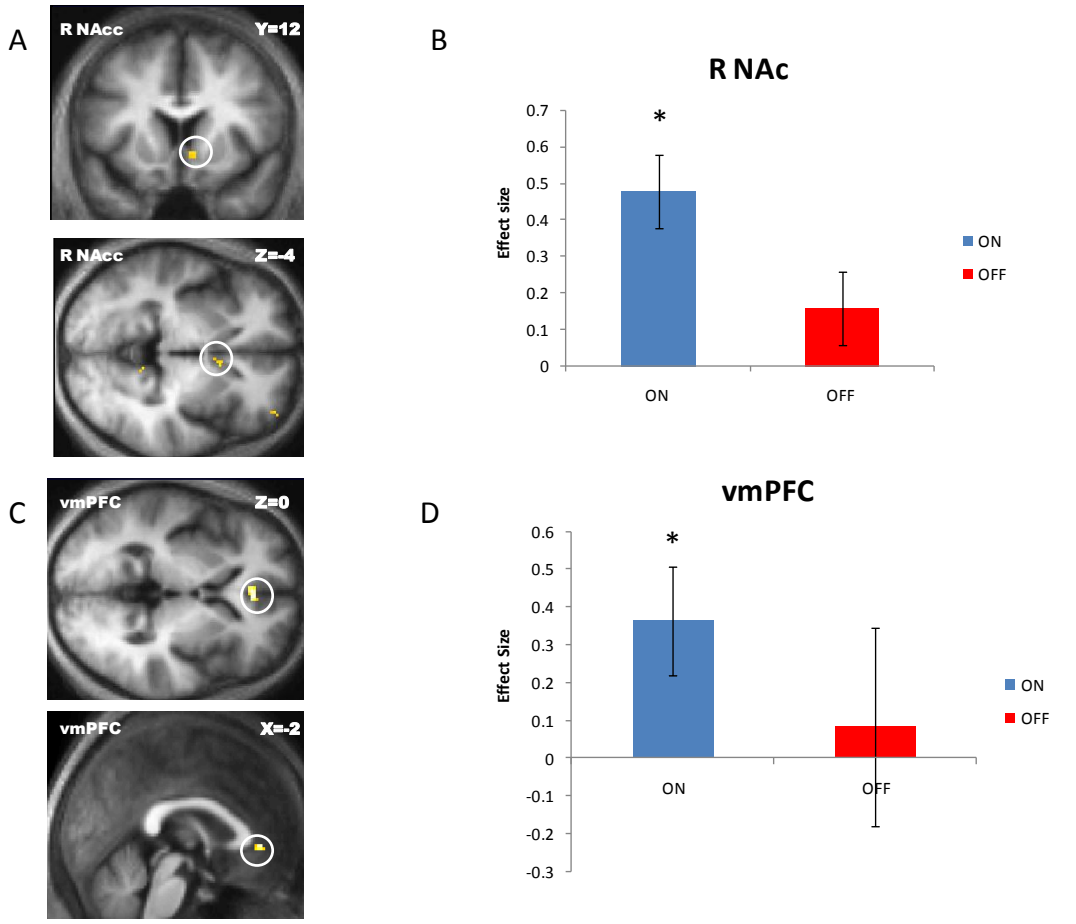

Figure IV.III: Brain activity correlating with the value of the chosen cue during the performance phase*. (A) Brain activity in right NAc correlated with the value of the chosen cue. Analysis performed over all correct trials (both ON and OFF) in a context where novel parings are presented. (B) A differential analysis between drug states reveals that this correlation was selective to the ON state (p=0.01). (C) Brain activity in vmPFC also correlated with the value of the chosen cue. Whole brain analysis performed over all correct trials (both ON and OFF). (D) Similar to the activity in NAc, this correlation between BOLD in vmPFC and the value of the chosen cue was only evident in ON but not in OFF state (p=0.02). The error bars represent SEM. Thresholds in SPM images set to p < 0.005 uncorrected.*

Although there were no differences in performance accuracy in the *training pairs* between the different drug states during the performance phase we checked separately for value related neural activity for these pairs during the performance phase. Akin to these behavioral findings and to the neuroimaging findings from the acquisition phase we did not find a significant correlation between BOLD activity and the value of the chosen cue in the training pairs during performance in either of the drug states.

4.4 Discussion

Behaviourally I show an effect of dopamine replacement therapy (DRT) on the ability of PD patients to select the highest valued stimulus in a probabilistic reinforcement learning task. The medication status during the *acquisition* task phase did not impact on successful task learning and the critical factor in determining patients' ability to select the best stimuli is the drug state during the performance phase, after the stimulus values have already been acquired successfully. Patients ON medication in the performance phase were significantly better at picking the best stimuli than when they were OFF medication irrelevant of the medication state they were in during the phase in which they acquired the stimulus contingencies. This effect was clearly seen in the novel parings and medication state did not impact performance when patients were required to select the best stimulus among pairs on which they had previously been trained.

A mechanistic basis for the behavioural findings is provided by my fMRI data, that specifically addressed the neural representation of stimulus value during the performance phase. Even when subjects had learned stimuli OFF DRT, and only given their DRT after learning had occurred, activity in nucleus accumbens and vmPFC encoded the value of the chosen stimulus during the performance phase. This suggests that, in contrast to previous accounts (Bayer & Glimcher 2005; O'Doherty et al 2003; Pessiglione et al 2006; Schultz et al 1997), reduced dopamine availability during learning does not impair value acquisition. In keeping with this, I did not find any behavioural or neural differences between the different drug states

in the acquisition phase of the task. Instead, the data show that decreased dopamine during performance resulted in an impoverished neural representation of stimulus value. It is of interest that the two structures highlighted in the data, the nucleus accumbens and vmPFC, are strongly associated with various forms of value prediction and prediction errors in reinforcement learning contexts (Dreher et al 2007; Luk & Wallis 2009; Matsumoto et al 2003). The pattern of findings observed, whereby stimulus value correlated with activity in these two regions in the ON state, implies that these brain areas can successfully represent the reward value of cues when patients are ON medication enabling successful performance for novel pairings. However, when this signal is degraded as seen in the OFF state, performance is impaired.

The involvement of the NAc during successful performance is particularly notable, since this structure is well known to control the immediate effects of dopamine on numerous aspects of performance (Berridge 2009; Berridge & Robinson 1998; Ikemoto & Panksepp 1999; Lex & Hauber 2010). The NAc is a site where the predicted value of stimuli are transformed into preparatory Pavlovian responses under a modulatory influence of dopamine (Berridge & Robinson 1998). We suggest that a preparatory response of approach is likely to be a key substrate for the behavioural patterns we observed in our task (Dayan et al 2006). This provides another reminder of the complexities inherent in a single neuromodulator (dopamine) supporting two apparently independent roles, namely reporting on appetitive prediction errors and influencing vigour (Boureau & Dayan 2011; Cools et al 2011; Ikemoto & Panksepp 1999; Niv et al 2007) .

A further important finding is the engagement of vmPFC in a context in which subjects made the correct choice between novel pairings of stimuli in the ON state, but not when subjects made incorrect choice in the ON state. This region is strongly implicated in valuation (Boorman et al 2009; Fitzgerald et al 2010; Gottfried et al 2003; Kable & Glimcher 2009; Plassmann et al 2010; Seymour & McClure 2008) across a range of experimental manipulations, with mounting evidence pointing to a specific role when subjects have to choose between distinct options with different values (FitzGerald et al 2009; Padoa-Schioppa & Assad 2006; Wunderlich et al 2010). This fits neatly with our observation that this region was engaged when subjects generated correct choices based upon an assessment of a learnt value difference between novel pairings. However, the data are intriguing in suggesting that the integrity of a dopamine input to this region is important for this form of value based decision. I acknowledge that I cannot be certain as to its precise role but two possibilities are immediately apparent, either dopamine is necessary for a stable value representation that can support generalisation or alternatively dopamine is necessary for a differencing operation needed when a subject, in the process of making a decision, needs to compare the value of distinct stimuli. I was unable to dissociate whether the neural value correlates were precursors to choice (action values) or the output of the choice process (chosen values) (Wunderlich et al 2009). It remains an open question for future research as to whether the deficit is due to a misrepresentation of pre-choice values and therefore due to a misrepresentation of values that are fed into a decision comparator, or a problem at the value comparison stage itself. Our study involved testing Parkinson's disease

patients, which although provides the best human model of dopamine depletion, carries the problem of whether the observations in patients can be generalised to the healthy population. Despite this caveat, the findings do lend support to the hypotheses (Berridge 2007; Berridge & Robinson 1998) and animal studies (Cannon & Palmiter 2003; Robinson et al 2005) which stress a major role for dopamine outside of learning.

In this study I was able to separate out the effects on dopamine on learning from the effects on performance and we found that the main effect of dopamine replacement therapy appears to be on the performance aspect of the task. At the neural level the improved performance in the ON medication state was associated with enhanced nucleus accumbens (NAc) and ventromedial prefrontal cortex (vmPFC) activity for the chosen cue value, an effect that was absent in the OFF medication state. The enhanced activity in the NAc and vmPFC in the ON medication state, which correlates with the improved behaviour, suggests that the enhanced cue value representation underlies the successful choice behaviour by allowing patients to select the best stimuli in novel contexts either by a more stable cue value representation or by improving the ability to compare the values of stimuli.

In summary, the improved performance in patients ON medication is not due to an effect of medication on learning and reflects a different effect of dopamine such as Pavlovian appetitive approach or motivational vigour associated with improved neural representation of cue value in a high dopamine state. By isolating the processes on which dopamine has the greatest impact, my findings point to likely

mechanisms that underlie common behavioural deficits seen in PD patients, both clinically and in various laboratory tasks, as well as providing a basis for future cognitive oriented therapies as well as shedding light on the fundamental role played by dopamine in reinforcement learning.

Chapter 5

The effect of valence on movement: a study of bradykinesia in Parkinson's disease

The experiment described in this chapter focuses on the motor deficits observed in PD. The aim of this study was to test whether movement speed in Parkinson's disease (PD) patients can be modulated by the specific nature of the motivational salience of possible action-outcomes, putatively demonstrating a link between the motor and cognitive deficits observed in PD.

5.1 Introduction

Akinesia is a cardinal feature of Parkinson's disease (PD)(Edwards 2008), consisting bradykinesia (slowness in executing movements), poverty of movement and a decrement in the size of repeated movements. Despite its impact, the precise cause of bradykinesia remains the subject of debate, with no single hypothesis providing a fully comprehensive account (Berardelli et al 1986; Hallett & Khoshbin 1980; Montgomery & Nuessen 1990; Sheridan & Flowers 1990; Sheridan et al 1987; Teasdale et al 1990).

Recent empirical findings and theoretical accounts suggest that bradykinesia, rather than being simply a manifestation of motor slowness, might reflect a specific deficit in the operation of motivational vigour in the striatum (Mazzoni et al 2007; Niv et al 2007; Niv & Rivlin-Etzion 2007). For example, compared with controls, PD patients could achieve similar speeds and accuracy of reaching movements, but did so more rarely, putatively demonstrating an implicit 'reluctance' to move fast (Mazzoni et al 2007).

A speeding effect of dopamine on action in response to rewards has been widely described (Moustafa et al 2008; Niv et al 2007; Salamone & Correa 2002). However, the effect of dopamine depletion on punishment avoidance is much less well understood and has not been formally tested in humans. One of the striking clinical characteristics of bradykinesia in PD is its variability (Blin et al 1990; Sheridan et al 1987), with the same patient being able to achieve very different movement times in different contexts. An extreme manifestation of this variability is *"kinesia*

paradoxical" where patients are suddenly able to move at near normal speeds, which usually occurs only in extreme aversive contexts (Critchley 1929; Rahman et al 2008). This class of observation motivated me to examine if winnable rewards and avoidable punishments might have differential effects on movement time. Furthermore, a differential effect would shed light on valence and vigour opponency between dopamine and its putative opponent (Boureau & Dayan 2011; Cools et al 2011; Daw et al 2002).

The use of rewards and punishments furnished me with an opportunity to test whether there is an effect of dopamine depletion, as manifest in the Parkinsonian state, on an ability to maintain a response plan or working memory trace in the face of distraction and whether this is valence specific. This in principle could explain some of the conflicting findings in the literature: PD patients are impaired when required to 'multitask' motor and cognitive tasks (Hausdorff et al 2003; Praamstra et al 1998; Shohamy et al 2006), although when working memory is explicitly tested, dopamine depletion reduces distractibility (Cools et al 2010; Crofts et al 2001). However in these tasks outcome valence was not explicitly manipulated, leaving unresolved the question of whether an impact of distraction may be context (valence) sensitive.

I developed a novel movement time paradigm involving winnable rewards and avoidable electric shocks, and tested PD patients and matched controls. Critically, I assessed movement time and not reaction times. The motivation here was to remove any confound of cueing, given the known sensitivity of PD patients to visual and auditory cues (Brown & Marsden 1988; Lewis et al 2000). Additionally, I was

specifically interested in measuring the time it takes to *execute* as opposed to *initiate* a movement, thereby focusing on the core deficit found in bradykinetic patients. In my paradigm, the faster the subjects performed an action the more likely they were to win money (in appetitive blocks) or to avoid an electric shock (in aversive blocks). I compared patients when OFF dopaminergic medication with controls. This means I tested patients in a more natural disease state, minimising as far as possible the effect of medication fluctuations and dose variations.

5.2 Materials and methods

5.2.1 Subjects

Twenty three adults (12 PD patients and 11 healthy control subjects) participated in the study, whose procedures were approved by the National Research Ethics Service, Moorfields & Whittington Research Ethics committee. Patients were recruited from the movement disorder clinic at the National Hospital for Neurology and Neurosurgery (NHNN). Control subjects were recruited through advertisements in public libraries or were spouses of patients. Written informed consent was obtained from all subjects and transport costs associated with participation were reimbursed. The participants were paid an extra fee of between £5 and £15 which was dependent on task performance.

Subjects were screened for psychiatric and neurological co-morbiditiy as well as current and past medication. They were also examined by a clinician and asked to complete several questionnaires, including a depression scale, an impulse control disorder screening questionnaire, and a mini-mental state examination.

5.2.1.1 Subjects with Parkinson's disease

Twelve English speaking early- to moderate-stage [H+Y stage- mean (SE) 2.4 (0.14)] (Hoehn & Yahr 1967) PD patients (eight males) aged between 48 and 82 years [mean (SE) 66.6 (2.6) years] participated in, and completed, the study. Eleven of the subjects were right-handed. They had on average (SE) 13.25 (0.66) years of education. Initial diagnosis of Parkinson's disease varied from 3 to 9 years [mean (SE) 5.45 (0.7) years]. There was no history of other major neurological or

psychiatric disease. Patients were on various regimens of anti-Parkinsonian medications; 11 subjects were taking carbidopa/levodopa combinations; one was receiving dopamine receptor agonists alone. Total daily dose of carbidopa/levodopa varied from 75/300mg to 250/1000 mg [mean (SE) 117/468 (19.6/78.7) mg]. (see table 5.5.1 for details of other medications) .

5.2.1.2 Control subjects

Eleven English speaking control subjects (six males) aged between 38 and 73 years [mean (SE) 61.72 (3.1) years], with no current major health problems or history of neurological or major psychiatric illness, participated in, and completed ,the study. Nine of the subjects were right-handed. They had on average (SE) 14.2 (0.8) years of education. Current medications included anti-hypertensive drugs (three subjects), lipid-regulating drugs (three subjects) and antidepressants (one subject) (table 5.5.1).

5.2.2 Experimental Paradigm

Both groups completed the computerized movement time task detailed below. Subsequently, subjects completed a battery of neuropsychological tests, comprising: (i) the Mini Mental State Examination to assess cognitive impairment (Folstein et al 1975); (ii) the Beck Depression Inventory (Beck et al 1961); and (iii) an impulse control disorder questionnaire **(**adapted from (Weintraub et al 2009)). In addition, the severity of clinical symptoms was assessed in the Parkinson's disease (PD) group according to the Hoehn and Yahr (Hoehn & Yahr 1967) five-point rating scale, and using the Unified Parkinson's Disease Rating Scale (UPDRS – all sections) (Fahn S 1987). Parkinson's disease subjects completed one test session in the

relative 'OFF' medication state, following a minimum of 12 hours withdrawal from all dopaminergic medication and omission of all slow release preparations for a minimum of 18 hours. The average Hoehn and Yahr rating for the patients was 2.4 [mean (SE) 2.4 (0.14)] and UPDRS was 48.5 [mean (SE) 48.5 (3.6)]. PD patients and controls were well matched for age ($F(1,21)=1.44$, $p=0.242$), education (years) ($F(1,21)=0.87$, $p=0.361$) and MMSE ($F(1,21) = 0.48$, $p=0.495$). PD patients had higher BDI and ICD scores however when compared to controls the differences only reached trend level significance (BDI ($F(1,21)=3.39$, $p=0.08$ and ICD ($F(1,21) = 3.05$, $p= 0.095$) (table 1)), we acknowledge however that this may be due to the small sample sizes. Parkinson's disease patients and controls were well matched in terms of age, education, sex, and on their neuropsycological test scores (see table 5.5.2). Control subjects also completed one test session.

5.2.2.1 Movement time task

Stimulus presentation and response recordings were conducted using Cogent software (www.vislab.ucl.ac.uk), programmed in Matlab (Natwick, MA). The task was designed to measure movement times in response to stimuli associated with rewarding or punishing outcomes. There were two types of trials: trials in which participants' aim was to win money and trials in which the aim was to avoid shocks. The task consisted of 6 interleaved blocks of 50 trials each, with blocks of 'money' trials alternating with blocks of 'shock' trials. The first block type was randomised between subjects.

Trials began with presentation of either a money or shock symbol for 2 seconds. The symbols were presented on a blue or yellow background, corresponding to

trials in which subjects could win money or trials where they should avoid shocks. This indicator of context was designed to remind participants of the current trial type. Background colours were counterbalanced across subjects. Participants were instructed to refrain from any action while the symbol remained on the screen. When the symbol disappeared, they were then required to press a key on the keyboard to start the trial. The trial would only start when the first key was pressed. Trials were self-paced. I opted for this design specifically to prevent the start of the trial being explicitly cued, in the light of the known effect of cueing in PD. (Brown & Marsden 1988; Lewis et al 2000) After commencing a trial, by pressing the first key, subjects then needed to press an adjacent key on the keyboard, using the same finger, in as quick a time as possible. On half of the trials (both in the money and shock trials), after the first key was pressed, a green flashing box appeared mid screen, which subjects were instructed to ignore. The role of this flashing box was to provide an attentional distractor. The flashing box remained on the screen until the trial was terminated by the second button press (see figure V.I for task depiction).

The time between the first and second button press was defined as the movement time. Following the second button press (i.e. completion of the trial), a screen was shown indicating trial outcome. In the 'money' blocks, participants either did or did not win 10p. In the 'shock' blocks, participants either avoided or received a shock. To incentivise fast movements, encouraging subjects to perform actions at maximal speed, and to reduce habituation, I varied the probability of outcomes (between 0.2 and 1 in an identical fashion for all subjects) such that receipt of reward or the

omission of punishment was linearly dependent on the speed of the associated movement time.

Subjects first performed a short practice session in order to familiarise themselves with the task. The instructions for the task were presented on the computer screen. During the practice session they neither received shocks nor won money.

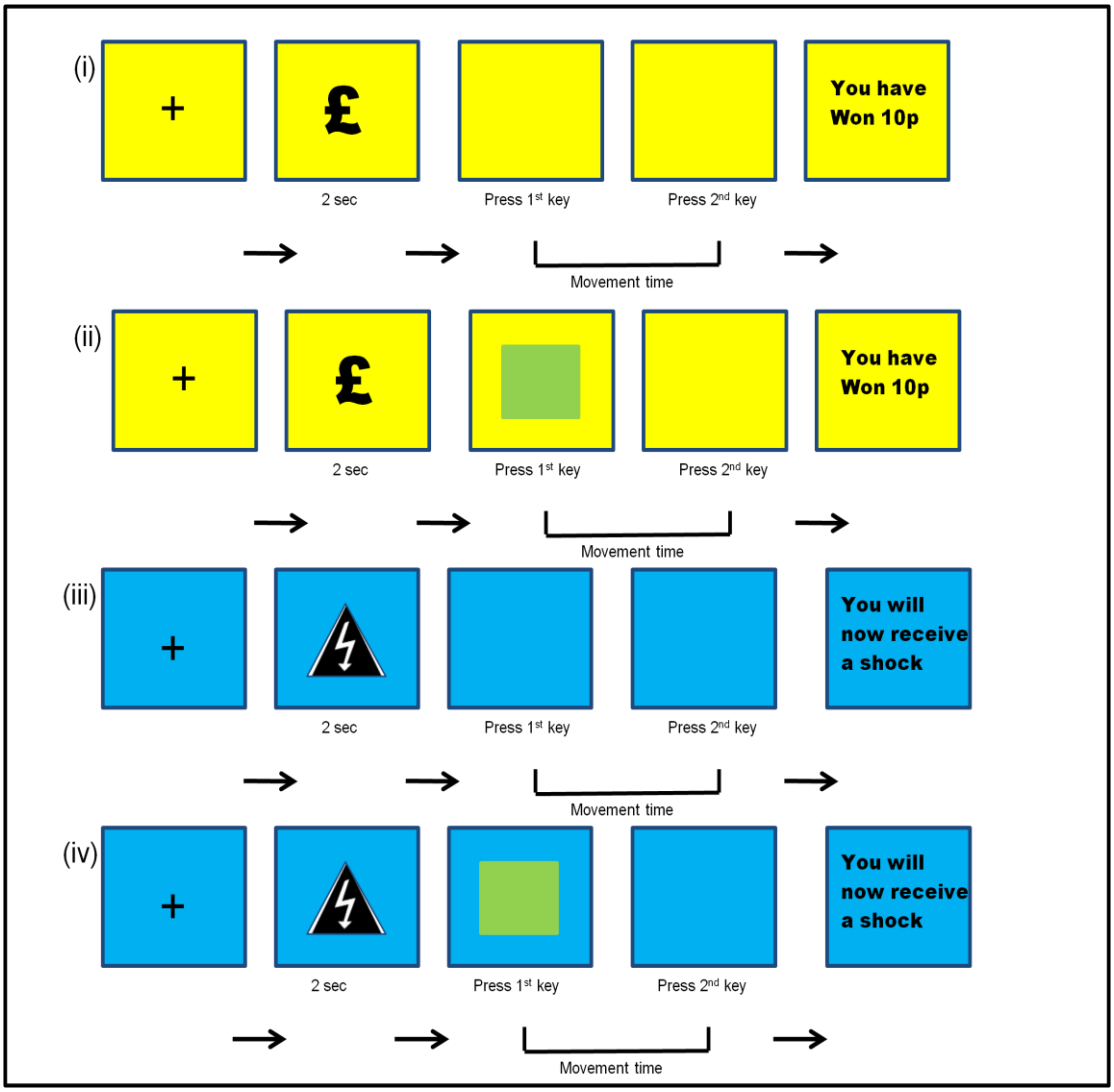

Figure V.I: Schematic of the movement time task

Trial types are illustrated as a function of outcome valence (yellow for money trials and blue for shock trials) and presence or absence of a distractor (green flashing square). There were two possible outcomes in the money trials; 'you have won 10p' or 'you have not won 10p' and there were two possible outcomes for the shock trials; 'you will now receive a shock' or 'you will now not receive a shock'. Subjects were exposed to 4 distinct trial types (see methods for details) comprising

(i) money trial without distractor,
(ii) money trial with distractor,
(iii) shock trial without distractor,
(iv) shock trial with distractor.

Failure to complete a trial correctly, for example by pressing the same button twice in error, resulted in no outcome being delivered (i.e. no money or shock outcomes). Although it is conceivable that subjects could have used this as an 'escape route' from aversive outcomes, when I tested this possibility explicitly I found that subjects failed to carry out trials only on very few occasions during testing [mean (SD), 2.47 (4.35) from a total of 150 trials]. Failure to respond on one trial did not impact on movement time on subsequent trials and mean movement time before, and after, this contingency was utilised did not differ significantly (TTEST $p>0.2$).

5.2.2.2 Apparatus

Participants were seated in a well-lit room in front of a desktop computer with a normal keyboard.

Electric skin stimulation

Two Digitimer boxes were fitted with circular electrodes. The triggers for the shock box were sent via the parallel port to the input on the shock box. Before commencing the task, participants had an electrode attached to the back of their non-dominant hand. They then underwent a shock titration procedure. This consisted of first establishing a maximal threshold level at which the electrical current was rated as very uncomfortable. Then, an automated staircase procedure was used to determine the level of shock for each individual that was 60% of their own maximal threshold.

5.2.3 Data Analysis

I initially focused on the overall effects of disease on movement time in the task, examining the differences in performance in the money versus the shock trials, and comparing the effects of these outcomes with those obtained in the control group.

I also examined effects of previous trials' outcomes on the movement times of subsequent trials by performing multiple regression analysis. Here I modelled separately the modulatory effects of receiving money compared with not receiving money on the previous trial; and the effects of receiving shock compared with not receiving a shock on the previous trial. I also included terms for the overall average effect on movement time of money and shock trials, anticipating that these would be different. I estimated the betas from the regression model and performed one sample t-tests on these at the group level to make inferences about the effect size of four factors:

$$MT=\beta_1 \times Money + \beta_2 \times Shock + \beta_3 \times M_{(t-1)} + \beta_4 \times S_{(t-1)} + \beta_5 \times D_{(t)} + \varepsilon$$

Where:

MT = movement time

Money = indicator variable for all money reward trials

Shock = indicator variable for all shock punishment trials

$M_{(t-1)}$= indicator (1/-1) of outcome of previous money trial

$S_{(t-1)}$ = indicator (1/-1) of outcome of previous shock trial

$D_{(t)}$ = indicator (1/-1) of whether distractor present.

ε = error term

All terms were entered simultaneously into the regression (without being orthogonalised).

Additionally, I performed an ANOVA examining the effect of the distractor on movement times, testing for a 3-way interaction between block type (money/shock), distractor (present/not present) and group (controls/patients). I also looked at the time taken from the appearance of the money or shock symbol until the first button press. This was in order to confirm that there was no difference in movement time between the two groups or valence conditions which could indicate differences in motor preparation times.

I excluded 2 sessions (one money session in a control subject and one shock session in a patient) where movement times in the first block were over 150% longer than the movement times in subsequent blocks for the same type of trial. I believe this incongruous performance in these subjects reflects an initial failure to understand the task demands which led to performance changing drastically between the first and subsequent blocks.

5.3 Results

My analysis indicated two main effects. First, I found an effect of group whereby patients were slower overall than controls $F(1,21)=15$, $p=0.001$. Second, I found an effect of valence such that both patients and controls were faster for shock compared to money trials (paired t-tests comparing money with shock trials in controls $T_{(1,10)} =2.51$, $p=0.03$; and in patients $T_{(1,11)}=3.49$, $p=0.005$) (figure V.II, table 5.5.3).

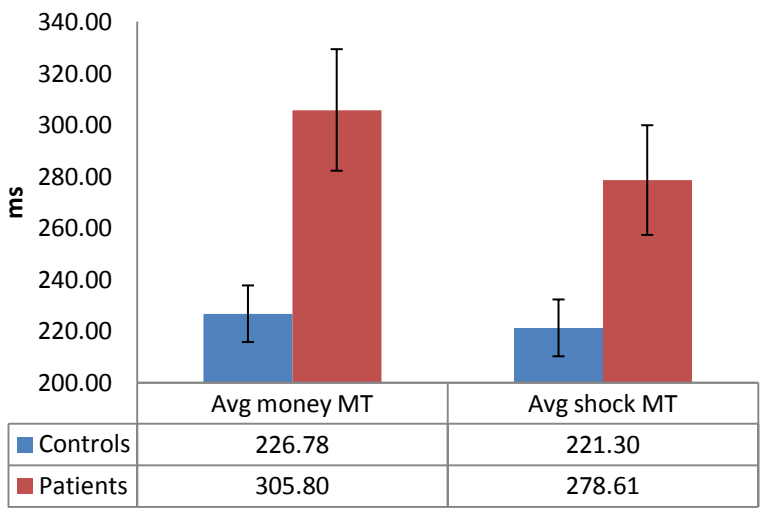

	Avg money MT	Avg shock MT
Controls	226.78	221.30
Patients	305.80	278.61

Figure V.II: Movement times raw numbers

Average movement times (ms) for money trials and shock trials for controls (blue bars) and patients (red bars). Error bars represent (two times) the standard error.

Crucially I observed an interaction between group (control/patients) and outcome valence condition (shock/money) $F_{(1,21)}=6.6$, p=0.017. The interaction was characterised by a bigger *difference* in movement time (MT) between money trials and shock trials in patients compared with controls (figure V.III).

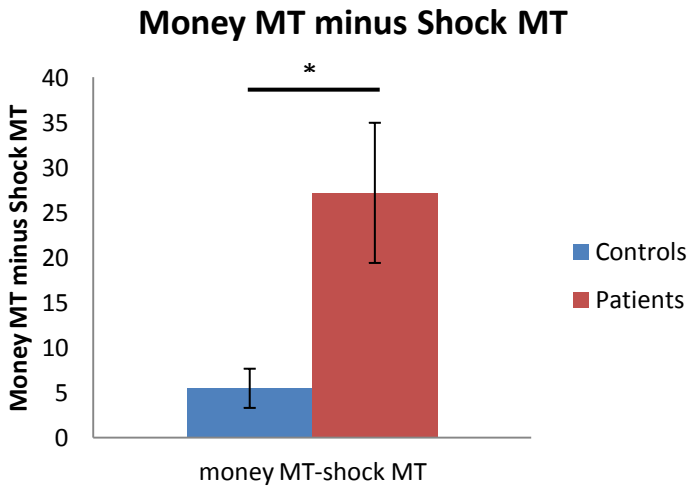

Figure V.III: Differences in movement time in money compared with shock trials

Movement time (ms) in money trials minus movement time in shock trials for controls (blue bars) and patients (red bars). Error bars represent (two times) the standard error of the difference.

I next examined the effect of outcome in a previous trial on movement time in the subsequent trial. I hypothesised that movement times would be influenced both by context (i.e. money compared with shock trials) and also experience on a previous trial, evident in a trial-by-trial sensitivity to rewards and punishments. For example, I expected that failure to achieve the desired outcome (i.e. not winning money or actually receiving a shock) on a previous trial would lead to faster movement on the subsequent trial. This is exactly what I found for the control group in the case that they failed to win money ($T_{(1,10)}$=-2.23, p =0.049 two tailed) . However, this speeding effect was absent in patients ($T_{(1,11)=-}$1.23, p=0.242 two tailed) (p>0.25). Both patients and controls responded in the same manner to receipt of a shock by tending to improve their speeds in the trials following shocks. However, this speeding was not statistically significant.

To ensure that the faster responding for shock in the patient group could not be explained by a prolonged motor preparation time, I examined the time taken from the symbol appearance to the first button press. There was no significant difference in this initial period of time between the groups (patients /controls) $F_{(1,21)}$=0 (p=0.9997) or conditions (shock /money) $F_{(1,21)}$=0.49 (p=0.464) or an interaction between group and condition $F_{(1,21)}$=2.5 (p=0.138) ruling out this possibility .

A differential effect of distractor on movement time was evident. In the repeated measures ANOVA there was a significant 3 way interaction between block (money/shock), distractor (present/not present) and group (controls/patients) ($F_{(1,20)}$ = 7.54 p=0.012) characterised by patients' movement time slowing when performing a shock trial where a distractor was present (figure V.IV)

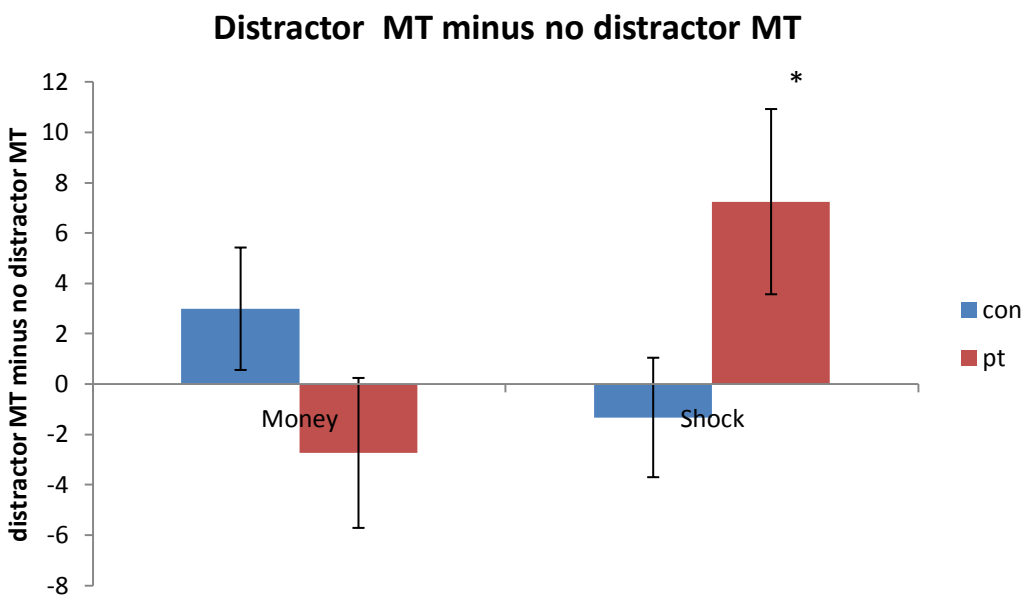

Figure V.IV: Effect of distractor on MT

Movement time (ms) in trials with a distractor present minus movement time with no distractor present, for money trials and shock trials, for controls (blue bars) and patients (red bars). Error bars represent (two times) the standard error of the difference.

The multiple regression analysis allowed me to look for more subtle differences in the modulatory effect of previous trials on movement time in the two groups while controlling for other factors. This confirmed my findings that not winning money in a previous trial had a significant effect on movement time in the control group (one sample t-test $p<0.001$ two tailed. $\beta_{3\ (\text{controls: effect of loss at t-1})}$: mean (SE) 3.86 (0.7)), showing that controls sped up significantly on a trial after they failed to win money. This speeding effect was absent in patients (one sample $p>0.4$, two tailed $\beta_{3\ (\text{patients: effect of loss at t-1})}$: mean (SE) 2.15 (2.7)).

5.4 Discussion

The most notable result is a valence asymmetry in the movement time of PD patients. This comparative failure to speed up in order to win rewards is consistent with previous findings in PD patients OFF medication (Moustafa et al 2008) and supports proposals that tonic dopamine levels control the rate and vigour of movements, possibly by signalling the average reward rate in the environment (Niv 2007; Niv et al 2007). This notion has been linked to the idea of impaired 'motor motivation' in PD, whereby there is a shift in the cost/benefit ratio of moving fast (Mazzoni et al 2007). Crucially, I find that although the response to rewards appears impaired in the PD group, the trial-by-trial response to punishments is not similarly impacted, a fact which has not previously been demonstrated. This finding highlights that in PD, dopamine depletion has a lesser impact on responses to punishments compared to rewards, and hints at a more complex role for dopamine in active avoidance. A critical aspect to the task is that I examined the effect of explicit contexts on movement time and compared subjects in dopamine depleted and non-dopamine depleted states. My findings indicate that bradykinesia is not simply related to movement, but rather to the way in which a hypodopaminergic striatum computes action values.

Importantly, I observed a difference in the effect of past monetary loss on subsequent actions in patients compared with controls, where subjects were given trial-by-trial feedback on whether their performance sufficed to merit a reward or avoid a punishment. If learning is effective, I expected a speeding up of movements following trials with negative outcomes (failure to win money or avoid a shock),

thereby improving the chances of achieving the desired outcomes on subsequent trials. This effect was clearly evident in the control group for rewards, but was absent in the PD patients. Patients did not speed up their movements after failing to win a reward despite physically being able to move faster, a fact they clearly demonstrated in the shock avoidance trials. The observation of failure to adjust movement time in the face of monetary loss in PD patients tallies with findings of impaired reward feedback learning in PD patients OFF medication (Czernecki et al 2002; Frank et al 2004). Of note, this trial-by-trial adaptation, whereby subjects speed up in response to a failure to win money has been observed previously albeit in the context of a probabilistic task in which this speeding was evident in both controls and patients (Moustafa et al 2008).

Finally, I found a detrimental effect of a distractor that was only evident in the shock trials in PD patients, indicating that here too there is an asymmetrical effect of valence. The context specificity of distraction has been demonstrated previously, with susceptibility to distraction being higher in PD patients when multitasking is required (Hausdorff et al 2003; Praamstra et al 1998; Shohamy et al 2006) but lower in working memory tasks when OFF medication (Cools et al 2010; Crofts et al 2001). Here I show that distraction is also valence specific. Given that PD patients can improve both their motor speed and accuracy of their movements with increased attention (Baker et al 2007; Cunnington et al 1995), I hypothesised that the hypodopaminergic state in PD led to decreased attending to appetitive stimuli compared with aversive stimuli, improving motor performance at the cost of an increased sensitivity to distraction in the aversive trials.

In summary, I provide evidence that bradykinesia is not a fixed, context-independent, deficit. I link the cognitive and motor deficits associated with the PD hypodopaminergic state by demonstrating that bradykinetic movements are dependent on the valence frame in which movements are executed. Such modulation is apparent in *"kinesia paradoxical"*, where PD patients can suddenly move quickly in exceptional circumstances (Critchley 1929; Rahman et al 2008) or indeed (though less closely associated with the theory) when explicit visual or auditory cues are present (Suteerawattananon et al 2004). Here I showed this effect in a controlled environment with conventional cues whose motivational salience is internally rather than externally assessed. Additionally I demonstrate that distractors play an important role in performance in PD patients and that this effect is also valence specific. These data have clinical implications, potentially yielding new strategies to increase the effectiveness of rehabilitation treatments.

5.5 Appendix

Table 5.5.1: Medications

	Patients (n=12)	Controls (n=11)
Sinemet (carbidopa-levodopa)	11	0
Stalevo (carbidopa-levodopa-entacapone)	1	0
Ropinirole	4	0
Trihexiphenidyl	1	0
Selegiline	2	0
Pramipexole	1	0
Co-Q10	1	0
Clonazepam	1	0
Anti-hypertensives	2	4
Anti-depressants (SSRI/SNRI)	1	1
Warfarin	1	0
Terazosin	1	0
Thyroxine	1	0
Statin	1	3
Ceterizine	1	1
Aspirin	1	2

SSRI = selective serotonin reuptake inhibitor. SNRI = Serotonin–norepinephrine reuptake inhibitors

Table 5.5.2: Neuropsychological data sets

	Patients (n=12)	Controls (n=11)
Age	66.6 (2.6)	61.7 (3.1)
Education (years)	13.2 (0.6)	14.2 (0.8)
MMSE	28.5 (0.3)	28.9 (0.4)
BDI	10.2 (1.5)	6.1 (1.5)
ICD	2.25 (0.7)	0.63 (0.54)

Values represent mean (SE). BDI = Beck Depression Inventory; MMSE=Mini Mental State Examination; ICD = impulse control disorder questionnaire.

Table 5.5.3: Raw data

	Patients (n=12)	Controls (n=11)
Mean MT money trial	305.8 (23.5)	226.7 (10.9)
Mean MT shock trial	278.6 (21.2)	221.2 (11)
Mean MT distractor *money trial*	304.4 (23.2)	228.2 (11.2)
Mean MT no distractor *money trial*	307.1 (23.2)	225.2 (10.8)
Mean MT distractor *shock trial*	277.2 (21.7)	220.4 (10.4)
Mean MT no distractor *shock trial*	269.9 (23.3)	221.8 (11.6)

Values are in milliseconds and represent mean (SE).

Chapter 6

Expectations and violations: Probing the role of dopamine in set shifting

The experiment described in this chapter focuses on the role of dopamine in set shifting. I used a pharmacological manipulation in healthy individuals to isolate the role of dopamine in set shifting while controlling for executive aspects of motor vigour and for responses to non-specific violations of sensory predictions, with the hope of better understanding the neurobiology underlying pathological behaviours associated with the hyperdopaminergic state.

6.1 Introduction

The aim of this study was to characterise the role of dopamine in set switching. Specifically, I measured the effects of (L-dopa) manipulation of dopaminergic neurotransmission on behavioural and neurophysiological responses to cues calling for a change in response set.

In brief, subjects were required to switch sets between a go and a no-go response, when they encountered an unexpected outcome following sequential presentations of the same target stimulus. Alternating between go and no-go sets enabled me to average over behavioural and physiological responses that did and did not involve motor activity and thereby focus on set switching per se, independent of action. Furthermore, by comparing responses to unexpected losses with unexpected null outcomes, I was able to study the role of dopamine in modulating responses to cues with (negative) valence, as opposed to a non-specific violations of sensory expectations.

The motivation for this work rests on the observation that dopamine may be essential for high-level set switching and action selection, as evidenced by studies in normal subjects (Mehta et al 2004) and Parkinson's disease (Cools 2006; Cools et al 2001a). In this setting, I was interested in how the neuromodulatory effects of dopamine may be implicated in set switching and the maintenance of an appropriate representation of contingencies in working memory. This is in contradistinction to the role of dopamine in signalling rewards or outcomes per se.

To characterise set switching behaviourally, I measured the number of correct responses following an outcome that indicated a reversal of contingencies (i.e. a set shift). In terms of the physiological responses underpinning set switching itself, I used fMRI to measure the responses to (surprising) outcomes at the point of reversal. These responses were quantified in relation to the immediately preceding trial, in which predicted expectations were fulfilled. In short, I used the neurophysiological response to violations (surprising or unexpected losses) to measure the neuronal activity responsible for a switch in response set. I then examined the effect of perturbing dopaminergic neurotransmission with L-dopa on these behavioural and physiological responses.

My hypothesis was that the L-dopa would impair set switching and, behaviourally, decrease the proportion of correct responses on the trial following reversal. Specifically, I predicted that this decrease would be greater when avoiding losses, as opposed to avoiding null outcomes in similar fashion to that which has been previously been found in PD patients ON dopaminergic medication (Cools et al 2006). Physiologically, I predicted that during sequential cued responses (in both go and no-go contexts), the succession of cues and contingent responses would be encoded by delay period activity in the prefrontal cortex. The itinerant dynamics of these high-level central pattern generators are selected and maintained by dopaminergic gating of cortico-striatal interactions (Frank & Claus 2006; McNab & Klingberg 2008). When expectations about outcomes are violated, this itinerant (attractor) activity is destroyed and a new (metastable) dynamical representation emerges (Friston 1997; Oullier & Kelso 2006; Rabinovich et al 2008). Dopamine

plays an important modulatory role in working memory, specifically via D1 receptors, and is thought to be central in maintaining robustness of working memory representations making them more resistant to distractors (Durstewitz et al 2000). We assume that in our study, the switching between metastable attractors rests upon a de-modulation of ongoing attractor dynamics by dopamine, when expectations are violated. In other words, an unexpected outcome (that signifies a change in contingencies) causes a reduction in mesocortical dopaminergic modulation of prefrontal activity, allowing for the emergence of a new pattern of firing and consequent set shifting. In summary, I predicted a reduction in prefrontal responses to unexpected outcomes (losses) that reflects a suppression of itinerant (working memory) activity which is known encoding contingencies that are no longer consistent with sensory input.

The functional anatomy of these effects should, I predicted, be expressed in the projection fields of the ascending dopaminergic projections from the substantia nigra and ventral tegmental area: namely, the striatum (nigrostriatal pathway), the medial prefrontal cortex (mesocortical pathway) and nucleus accumbens (mesolimbic pathway) (Arias-Carrion & Poppel 2007; Haber 2003; Robbins 2000). My primary hypotheses concerned reductions in prefrontal responses at the point of reversal (i.e., to surprising or unexpected outcomes). However, I hoped to see similar or reciprocal subcortical responses.

I therefore first tested for a violation effect throughout the brain, in the hope of identifying significant responses in these regions (by comparing trials with unexpected outcomes with the preceding fully predicted trial), while using a fully

balanced design to control for any effects of executive aspects of motor vigour and non-specific sensory surprise on these representations. I then tested for the effects of dopamine within these regions, anticipating that violation-dependent effects would be attenuated under L-dopa. This is largely what I found; however, to my surprise L-dopa actually reversed the violation or surprise-dependent decreases in prefrontal responses.

6.2 Materials and Methods

The study and its procedures were approved by the UCL Research Ethics Committee.

6.2.1 Participants

All participants gave written informed consent. Only male participants were included to avoid menstrual cycle-dependent interactions between gonadal steroids and the dopaminergic system (Becker & Cha 1989; Dreher et al 2007). Sixteen men [age mean (SE) 23.8 (1.65)] completed the study. Two subjects (not included in the above analysis), one in the dopamine session and one in the placebo session were excluded due to excessive drowsiness leading to very low response rates during the scanning session.

15 of the subjects were right-handed and one was left handed. All were fluent English speakers with no history of other major neurological or psychiatric disease and no concurrent medication use.

6.2.2 Stimuli and task

I used a novel set switching task (see task depiction). Stimuli consisted of one of 4 Hiragana symbols presented in white fonts on a black background. All trials followed a similar sequence. Initially, a stimulus was shown on-screen for 200msec, before being masked (by a composite of all 4 Hiragana symbol characters) for a further period of time determined by the subject by subject reaction time (RT) measurement (see below). Subjects were required to make an appropriate response following stimulus presentation, either gripping ('Go') or omitting a grip

response ('NoGo'). After the individualised RT had elapsed, a coloured border appeared (500msec) indicating the type of response made on that trial. If subjects gripped, a yellow border appeared around the stimulus and if they did not grip a blue border appeared. This was to provide feedback to the participants so they could assess whether their intended responses were performed adequately within the allocated time (so that they know to adjust their behaviour if their gripping is not fast/strong enough). Following the feedback screen subjects were presented with an outcome screen (750 msec) which was different depending on which type of block the subjects were performing. In the 'null' block subjects could either continue to receive or not receive one pound (represented by a pound coin or an empty circle), in the 'avoid loss' block subjects could either not lose or lose one pound (represented by an empty circle or a pound coin picture with a cross through it), before commencement of the next trial.

Blocks of sequential trials were undertaken, lasting 17 min. There were two types of block, grouped according to the possible feedback subjects could receive following response execution. In block type 1 ('null' blocks), surprising set shifts were signalled by null outcomes (a blank circle). Conversely, in block type 2 ('avoid loss' blocks), surprising set shifts were signalled by salient loss outcomes (a cross overlying a pound coin). Correct responses in null blocks were signalled by a pound coin picture; in avoid loss blocks correct responses were indicated by a blank circle. These outcomes pertained to real monetary reward (see reward schedule below). Critically, because set shifts were rare events and the contingencies were entirely deterministic rather than stochastic, participants formed a strong expectation of

performing correctly and receiving expected feedback. In addition, by using fully deterministic outcomes we ensured that minimal learning was required to successfully alter behaviour, thereby allowing us to focus on the effect of a set shift and not on learning effects. Although it would seem that this method does not easily allow for separation of the set shift from the outcome, the behavioural trial of interest was actually the trial *after* the unexpected outcome thereby eliminating this problem. In the imaging data we focussed on the two trials preceding the switch and were specifically interested in examining the neural responses to fully expected outcomes and contrasting them with fully unexpected outcomes. In the imaging data we did not in fact examine the trial in which behaviour changed, only the brain responses which preceded the behavioural change. Moreover, running null and avoid loss contingencies in separate blocks entrained vigorous prepotent responses and avoided rapid contextual changes. 6 blocks were run in total (3 'null' 3 'avoid loss' blocks).

Occasionally, contingencies for the pair of stimuli in each block would switch with each other ('set switch'). On these trials a stimulus for which subjects previously had to grip (Go response) became a stimulus for which they had to not grip and vice-versa. To ensure that set shifts were unpredictable and that subjects had fully understood the contingencies before a set shift, I set a minimum constraint of 5 correctly executed trials (10 before the first set switch) per contingency mapping before a set switch was permitted. On trials subsequent to this, there was a 50% probability of switch per trial.

6.2.3 Procedure

Overview

We employed a within-subjects design, with each patient attending twice. All subjects performed the task in two different drug states, either on placebo (Cacit 1.25- contains 500mg calcium) or L-dopa (Madopar 187.5 mg contains 150mg Levodopa). This order was randomised, and sessions were scheduled a minimum of one week apart to ensure complete drug washout. Different Hiragana symbols were used for each visit (see figure VI.I for details of the task). The L-dopa and placebo were mixed with orange squash and both the participants and the investigator were blinded to the order of the drug/placebo. After ingestion, participants waited for 60 minutes to ensure maximum peak plasma drug concentration according to L-dopa pharmacokinetics (Khor & Hsu 2007). Participants performed both sessions at the same time of day to control for any diurnal fluctuations in baseline neurotransmitter levels.

Reaction time measurement

On the first visit, prior to L-dopa/placebo administration, baseline grip reaction times were measured to calibrate the subsequent set switching task. The task set-up was similar to the general trial sequence for the switch task. Two fractals were presented, one cueing a fast grip and one requiring omission of a grip response. These contingencies were explicitly described before commencing the task, and feedback was given following each response. The average +/- 2 standard deviations of RT (of grip trials), averaged over 15 presentations was used as the upper threshold time for a response in the set switch task. This controlled for intrinsic

within subject variability in response speed. Average individualised RT's (including 2 standard deviations of the mean were: [Mean (SE) 620ms (35.2)].

To familiarise subjects with the structure of the task they undertook a short practice block before scanning began in both sessions. During the practice session, subjects performed the identical task, except with different Hiragana symbols.

Payment schedule

On one of the days, after the scanning was completed, subjects underwent a structural scan. On the second session after task completion subjects received payment. To ensure incentive compatibility (i.e. so that subjects knew that each trial had the potential for real monetary loss or gain) 15 trials of each block type were randomly selected across sessions and paid out for real.

Set switching task

Each of 3 scanning runs lasted approximately 17 min, and consisted of 2 blocks ,one in which the aim was to avoid null outcomes and one in which the aim was to avoid monetary loss, in randomised order.

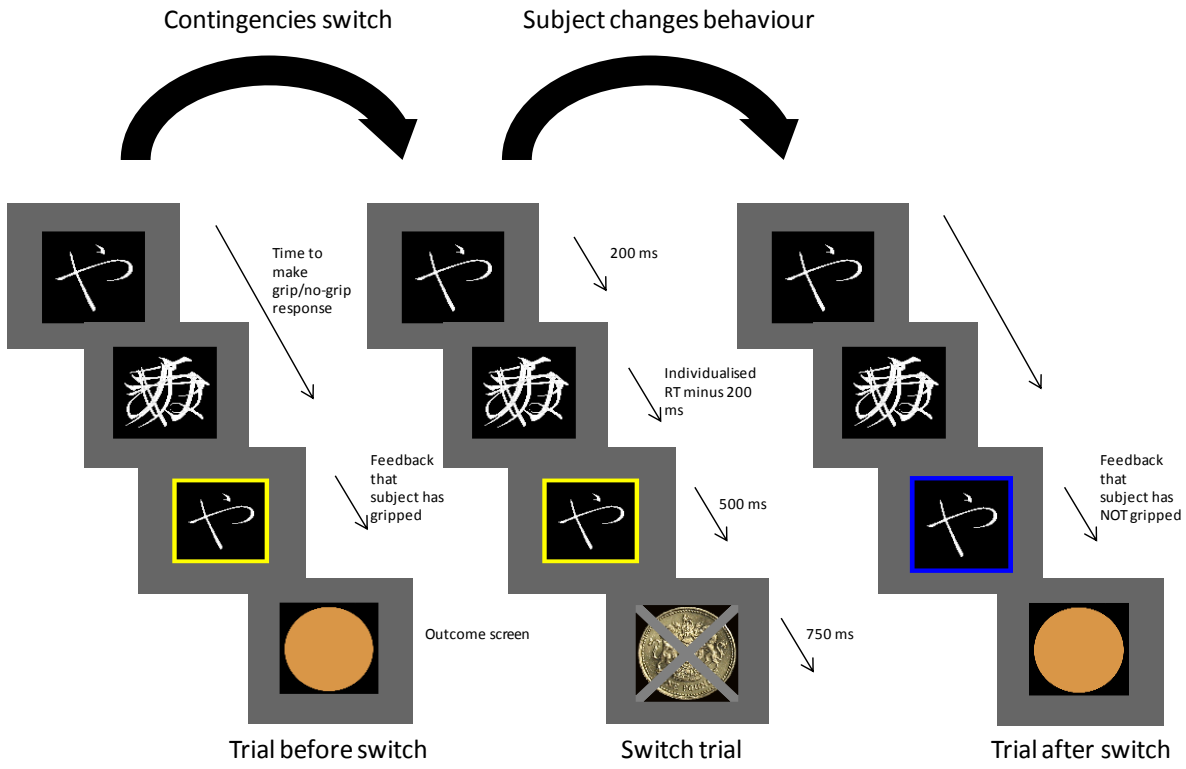

Figure VI.I : Task depiction. *Above is an example of a grip to avoid loss trial in the task. The subject grips in response to the Hiragana character and then receives feedback that the grip has been registered (indicated by a yellow border) and the outcome informing them they have not lost money (indicated by an empty circle). On the next trial the subject grips again however this is a switch trial and therefore when the subject grips they do not win money (indicated by a pound sign with a cross through it). On the next trial the subject switches their behaviour and now does not grip in response to the symbol (indicated by the blue border) and now does not lose money. The same process is occurring in parallel with a different symbol which switches from a no-grip to avoid loss to a grip to avoid loss at the same time. Also not depicted are the null outcome trials in which subjects can either continue to receive one pound (indicated by a pound coin) or receive a null outcome (indicated by an empty circle).*

MRI scanning

The study was conducted at the Wellcome Trust Center for Neuroimaging, at UCL using a 3T Siemens Allegra scanner equipped with a Siemens head coil. Anatomical images were acquired using magnetization-prepared rapid-acquisition gradient echo scans, which were followed by 1-mm-thick axial slices parallel to the anterior commissure–posterior commissure plane. Functional scans used a gradient echo sequence; repetition time, 2.86s; echo time 25 ms; flip angle 90 degree; matrix size 128x72; field of view 192 mm; slice thickness, 2 mm and interslice distance factor of 1mm. A total of 44 axial slices were sampled. The in-plane resolution was 3 x 3 mm. Functional imaging data were analyzed using statistical parametric mapping software (SPM8; Wellcome Trust Centre for Neuroimaging, London, UK; http://www.fil.ion.ucl.ac.uk/spm). Images were realigned with the first volume (after discarding the first six dummy volumes) and unwarped, normalized to a standard echo-planar imaging template based on the Montreal Neurological Institute reference brain, resampled to 3 x 3 x 3mm voxels, and spatially smoothed (8 mm full width at half-maximum).

6.2.4 Data analysis

Behaviour

The critical behavioural trial occurred on a switch. Here, subjects received surprising feedback, either a hedonically salient loss or neutral null outcome. This instigated a rapid alteration of behaviour following this violation of expectations. There were 4 types of trial –Go to avoid loss, NoGo to avoid loss, Go to avoid null, NoGo to avoid null and 2 drug states – placebo and L-dopa with a fully crossed

design. Accuracy (proportion correct responses) was computed for each. Comparisons between trial-types at the group level were performed by entering trial specific results into a three-way (drug x action x valence) repeated measures ANOVA.

Trials were separated into 'Go' and 'NoGo' trials for each of the valence outcomes (null /avoid loss) and the accuracy for the different drug sessions were computed separately. We also examined reaction times and gripper vigour differences. We checked late error Go versus NoGo trials in the separate groups, and performed a repeated measures ANOVA looking for interactions between the manipulated factors; null/avoid loss, Go/NoGo, L-dopa/placebo.

fMRI analysis: whole-brain general linear model parametric analysis.

Analysis of functional MRI (fMRI) data proceeded by a hierarchical analysis. At the within-subject level, a general linear model (GLM) was constructed containing regressors indicating each relevant trial type. There were 8 variables: Onset Go trial before switch, $T(Go)_{switch-1}$; Onset NoGo trial before switch, $T(NoGo)_{switch-1}$; Onset Go switch trial, $T(Go)_{switch}$; Onset NoGo switch trial, $T(NoGo)_{switch}$; Onset Go trial after switch, $T(Go)_{switch+1}$; Onset NoGo trial after switch, $T(NoGo)_{switch+1}$; Onset remaining Go trials; Onset remaining NoGo trials. The second model also contained 8 regressors: Onset correct trial before switch, $T(corr)_{switch-1}$; Onset incorrect trial before switch, $T(incorr)_{switch-1}$; Onset correct switch trial, $T(corr)_{switch}$; Onset incorrect switch trial, $T(incorr)_{switch}$; Onset correct trial after switch, $T(corr)_{switch+1}$; Onset incorrect trial after switch $T(incorr)_{switch+1}$; Onset remaining correct trials; Onset remaining incorrect trials. Trial-specific activations were modelled as box-car

functions, with durations set according to entire trial length on an individual subject basis, convolved with the canonical hemodynamic response function. Data from one subject had to be excluded from the second model due to a very high overall accuracy (>95%) thereby preventing out ability to contrast the correct with the incorrect trials.

Low frequency drifts were excluded with a high-pass filter (128-s cutoff). Short-term temporal autocorrelations were modelled using an AR(1) process. Motion correction regressors estimated from the realignment procedure were entered as covariates of no interest. Statistical significance was assessed using linear contrasts of the regression coefficients from the GLM, generating statistical parametric maps (SPM) of t values across the brain for each subject and contrast of interest. Placebo and dopamine sessions were separately analysed, and corresponding contrast images were taken to the second level as per a hierarchical random-effects analysis, entering contrasts for each subject in placebo and levodopa sessions into a group level GLM. Paried t-tests were used to make within subject comparisons of drug effects.

Anatomical localization was carried out by overlaying the t-maps on a normalized structural image averaged across subjects, and with reference to an anatomical atlas (Naidich 2009). All coordinates are reported in MNI space (Mazziotta et al 1995).

Region of interest analysis

Our primary regions of interest were the principal projection fields of midbrain dopaminergic afferents, specifically ventral and dorsal striatum and prefrontal

cortex (Haber 2003). After identifying significant voxels on a whole brain analysis, effect size data (beta values) were extracted from a 4mm sphere centred on peak activated voxel.

6.3 Results

Behaviour

To measure the efficacy of set switching, we calculated response accuracy (i.e. the proportion of correct responses) on trials subsequent to a switch of contingencies. Initially we characterised normal adaptive behaviour on placebo, where we found subjects were better at switching after a surprising/unexpected loss compared with an suprising/unexpected null event (paired t-test, null vs loss; T=-3.21 , p=0.006). This was entirely consistent with our prediction that loss would be a more potent catalyst for behavioural adaptation than null events.

We also found a switch type (loss/null) by drug (L-dopa/placebo) interaction ($p=0.004$, $F_{(1,15)}=11.73$) in the performance accuracy in the trial after the switch trial. As hypothesised, L-dopa obliterated the observed difference in performance following trials signalled by a hedonically salient surprising loss versus a surprising null event, such that accuracy on levodopa was now equivalent for both conditions (L-dopa: avoid vs null, paired t-test; T= 1.0, p=0.332), (figure VI.II).

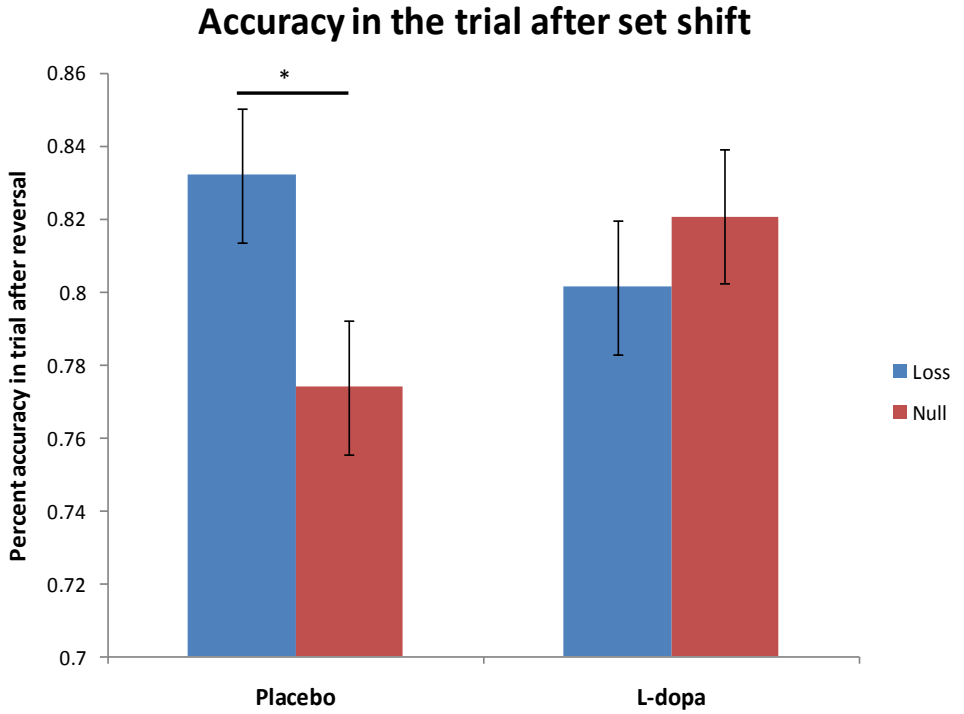

Figure VI.II: *Represented is the percent accuracy in the trial after set switch for the two different switch types separately (loss – blue bars/null-red bars) and for the two drug conditions (placebo and L-dopa). Error bars are within subject error bars.*

Critically, there was no main effect on accuracy of drug (L-dopa/placebo) ($F_{(1,15)}$ = 0.125, p=0.728) or of switch type (loss/null) ($F_{(1,15)}$=1.73, p=0.208) or action (Go/NoGo) (F=0.249, p=0.625). Thus accuracy was balanced, equivalent across each of these factors, such that we could specifically attribute behavioural changes following levodopa administration to a loss of the normal switching performance differential following hedonically surprising events. There was no interaction between switch type (loss/null) and action (Go/NoGo) ($F_{(1,15)}$=2.149, p=0.163), or drug (L-dopa/placebo) and action (Go/NoGo) ($F_{(1,15)}$=0.02, p=0.888).

fMRI

We initially measured baseline neural responses to a violation of expectations, given our central neural hypothesis that such responses drive subsequent behavioural set shifts. Thus, we computed the contrast [$T_{switch-1} – T_{switch}$], which localises differences in neural responses for trials with fully expected outcomes $T_{switch-1}$ compared with neural responses for surprising switch trials T_{switch} (where the outcome was unexpectedly incorrect) (figure VI.IIIA). This contrast between expected and unexpected responses was expressed in the vmPFC and NAc.

As with the behavioural data, we also found an interaction in the parameter estimates from the vmPFC (central coordinates x=9, y=26, z=-8) and a significant interaction between switch type (avoid/null) and drug state (placebo/L-dopa) ($F_{(1,15)}$=5.05, p=0.04) (figure VI.IIIB). Of note we also found a similar pattern of activation in the NAc (central coordinates x=-3, y=11, z=-2), whereby successful set shifting in the placebo avoid loss trials was associated with decreases in NAc activity with the reverse pattern in the L-dopa group. This effect only reached trend significance in the placebo group (one sample t-test placebo avoid T=-1.89, p= 0.077; one sample t-test L-dopa null T=-2.19,p= 0.044;) and there was no interaction.

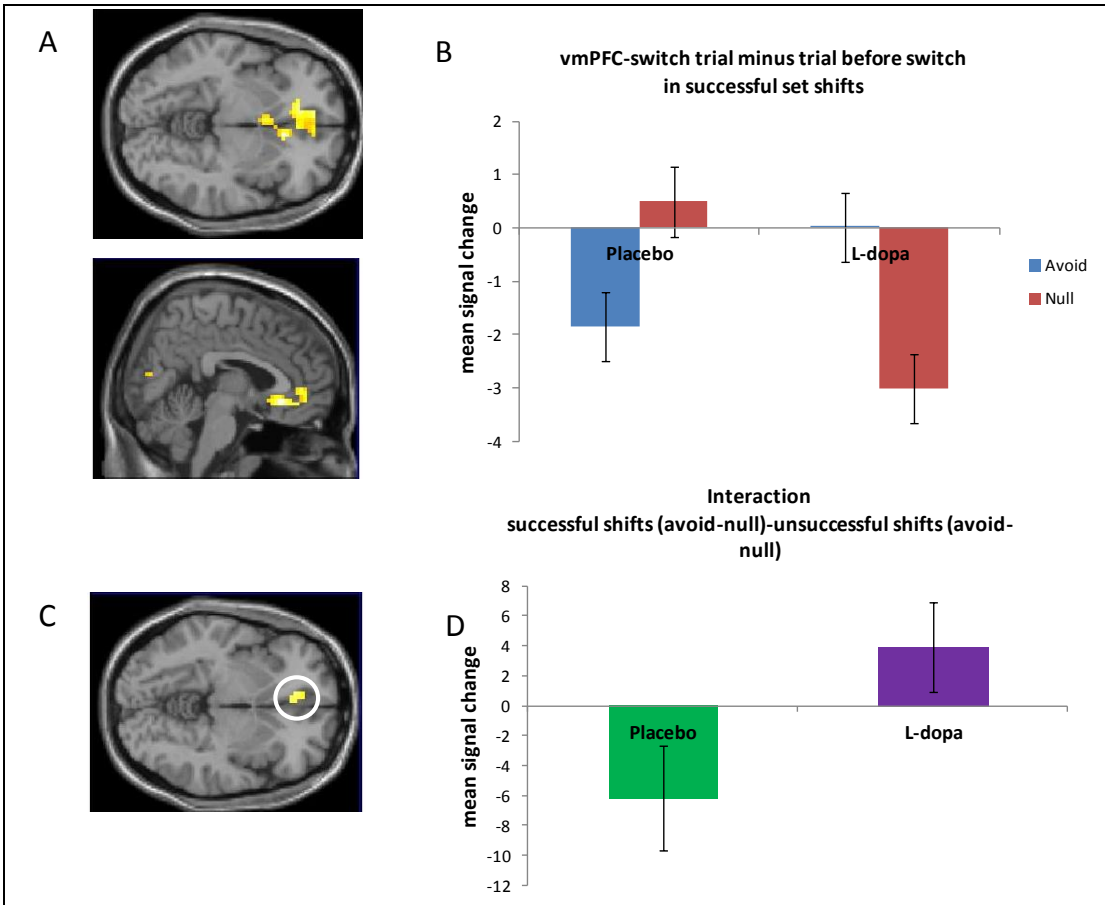

Figure VI.III : Neurophysiological responses to violations in expectations. *(A) Contrasted is the BOLD activation pattern in the trial before the switch (fully expected) minus the switch trial (unexpected) over all conditions (placebo/L-dopa and avoid/null). From the comparison of these two trial types activations are seen in the vmPFC and NAc. (B) Contrast estimates for the switch trial minus the trial before the switch, contingent upon subsequent successful set shifting (i.e. a behaviourally validated response). These show the difference in neural response between trials with an expected versus an unexpected outcome when a successful set shift occurs on the subsequent trial (for the different trial types (loss trials-blue bars, null trials-red bars). PFC neuronal activity is decreased during unexpected loss in the placebo group. In the L-dopa group this decrease is absent on the loss trials and is paradoxically only present in the null trials. Represented are the mean parameter*

estimates (beta values). Error bars are within subject error bars. (C) Contrasted is the interaction in the placebo group between successful and unsuccessful set shifts in the different trial types for avoid loss compared with null blocks (the contrast is successful shifts (avoid minus null) minus unsuccessful shifts (avoid minus null). (D) Beta plots representing the interaction in panel C. Illustrated is a decrease in activity in vmPFC in the behaviourally validated (successful) set shifts in the placebo group compared to the unsuccessful shifts, illustrating the link between the decrease in vmPFC activity on switch trials which result in successful behavioural shifting. Plotted are the within subject differences in the two drug conditions (placebo – green bar, L-dopa-purple bar) highlighting a decrease in activation in the placebo group when presented with an unexpected loss, an effect absent in the dopamine group. Plotted are the mean parameter estimates (beta values) and error bars represent standard error of the mean. The beta plots for the successful and unsuccessful trials only includes data from 15 subjects (see methods for details).

Basal ganglia loops are topographically distinct, however there is significant integration and overlap between the various loops and pathways (Haber 2003). Moreover, suppression of responses at one level of a neural hierarchy are often coupled to enhanced responses at complementary levels because of reciprocal information passing. Given this integration we anticipated that we would find reciprocity of response patterns between the mesocortical pathway with its projections to the vmPFC, and the nigrostriatal pathway with its projections to the caudate and putamen (Haber 2003).

The reverse contrast [$T_{switch} - T_{swich\text{-}1}$] (fully unexpected minus expected trial) highlighted increased activation in the caudate, insula and thalamus for trials with unexpected outcomes (figure VI.IV.A). Within these regions we asked whether L-dopa also modulated this pattern of enhanced activation. We indeed found that dopamine modulated caudate activity (central coordinates x=-12, y=5, z=7) in the unexpected minus expected trials (figure VI.IV.B). Contrast estimates demonstrate a clear reciprocity between the vmPFC and caudate activations (figure VI.IV.C).

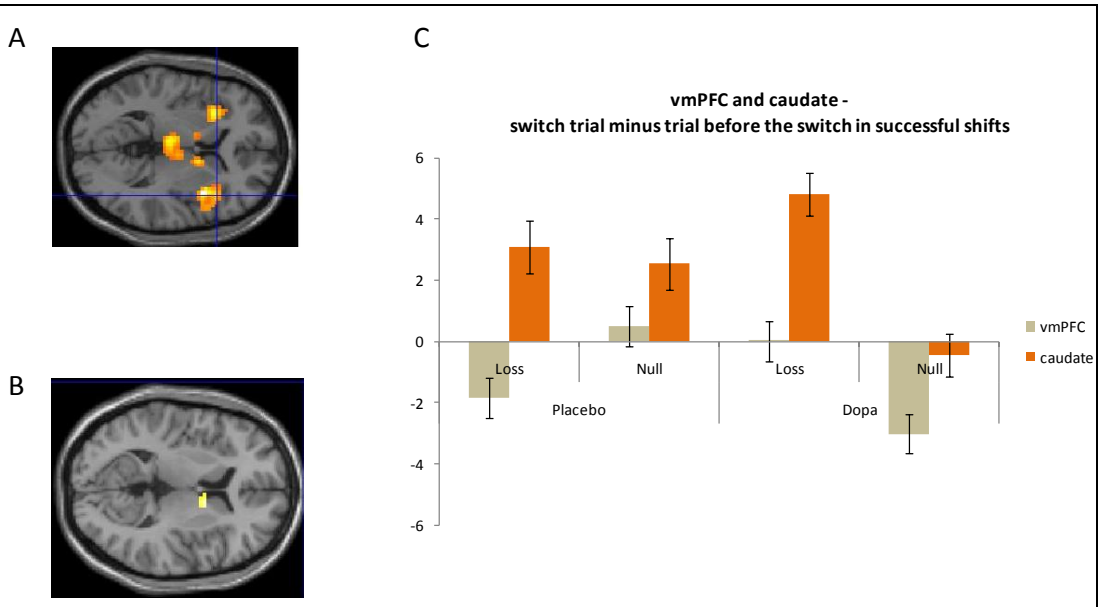

Figure VI.IV : Neurophysiological responses to violations in expectations.

(A) Contrasted is BOLD activation pattern in the switch trial (unexpected) minus the trial before the switch (fully expected) over all conditions (placebo/L-dopa and avoid/null). From the comparison of these two trial types activation in the insula, thalamus and caudate emerge. (B) Contrasted is the BOLD activation pattern in the switch trial (unexpected) minus the trial before the switch (fully expected) in the L-dopa minus placebo data. (C) Beta plots for switch trial minus the previous trial (before the switch) in the caudate and the vmPFC. The beta plots represent the difference between an unexpected and expected outcome when a successful set shift occurs on the subsequent trial (behaviourally validated set switch) for the different trial types. This illustrates an apparent dip in neuronal activity in the placebo group associated with an unexpected loss in the vmPFC. In the L-dopa group this effect is absent on loss trials and paradoxically is only present in the null trials. Conversely, the activity in the caudate shows an opposite pattern to that observed in the vmPFC. Represented are the mean parameter estimates (beta values) and error bars are within subject error bars. The beta plots for successful and unsuccessful trials only includes data from 15 subjects (see methods for details).

6.4 Discussion

Here, I aimed to characterise the influence of dopamine on the neurobiological processes that support set switching. I employed a stringent balanced design to control for purported possible effects of dopamine on movement and motor vigour (requiring both action and the omission of action as a response), and to disentangle effects of dopamine on non-specific violations of sensory predictions from salient, negatively hedonic surprise.

The appropriate shifts of behaviour in the face of aversive outcomes is a central evolutionary skill and the inability to correctly learn from punishments is a feature of many pathological lesions and conditions (Bechara 2005; Patterson & Newman 1993) . I aimed to elucidate the precise role of dopamine in this adaptive ability and as such to provide an explanation for impairments observed in disorders in which abnormal dopaminergic transmission is implicated (Berke & Hyman 2000; Evans et al 2009).

In the normal (placebo) state, subjects are highly skilled at rapidly altering behaviour following unexpected negative violations of expectations, but conversely are significantly less successful at set switching when the violation of expectations is signalled by a null outcome. Adaptive survival mechanisms, conferred by evolutionary selective pressures, are necessarily tuned to the avoidance of significant loss. It is therefore unsurprising to find such a striking asymmetry in behaviour determined by the hedonic valence of outcomes. Moreover, one would

expect the dopaminergic system to be optimised for such behaviour in healthy individuals.

This normal pattern of behaviour is in stark contrast to that observed when dopaminergic transmission is perturbed by the administration of L-dopa. On L-dopa, the ability of subjects to shift behaviour in response to negative violations of expectations was obliterated, with a significant decrement in the performance of subjects in this condition. Furthermore, when subjects were given L-dopa, we observed a reversal of the normal pattern of behaviour whereby subjects were better at set switching in response to null outcomes, albeit non significantly but leading to a two way drug by valence interaction. Simultaneously collected imaging data revealed a potential mechanism of this observed behaviour.

Dopamine plays a central role in the maintenance of working memory (Fuster 2001; Sawaguchi & Goldman-Rakic 1991; Stuss & Knight 2002b) and the representations of behavioural sets by delay period activity in the prefrontal cortex (Funahashi et al 1989; Fuster & Alexander 1971). Neurons in this region do not only code for stimuli and actions but also convey behavioural context (Asaad et al 2000). A current behavioural set is hypothesised to be encoded by a specific pattern of itinerant dynamics representing an entrained sequence of predicted states (Friston 1997; Rabinovich et al 2008). This entrainment is supported by dopamine, hence a strong prediction is that withdrawal of dopaminergic modulation in response to violations of expectations is required to allow the release from one pattern of attractor dynamics and the subsequent establishment of an alternate pattern thus facilitating a set switch. As predicted, we observed this pattern both in the vmPFC and NAc,

key projection structures of the dopaminergic pathways (Arias-Carrion & Poppel 2007; Haber 2003).

Significantly, this effect was marked only for hedonically surprising outcomes with negative valence, supporting the behavioural observation of enhanced performance after adaptively relevant losses. Moreover, disrupting brain dopamine levels by pre-treating with L-dopa not only attenuated the difference between negatively valenced and null events, but in fact reversed this pattern of neural activation.

The PFC also has been suggested to represent action (Fuster 2001), and dopamine has been shown to influence motor responses (Salamone et al 2003), it was therefore essential to control for motor execution when attempting to dissociate the role of dopamine in set shifting. Importantly, our observed effects, both subcortically and in vmPFC, were invariant to the execution or omission of a motor response. This strongly supports the notion that this network encodes a behavioural set per se, independent of the commission of action, and that dopamine enables the maintenance of a representation of specific state-sequences, rather than simply facilitating movement.

It is entirely understandable that perturbation of the dopaminergic system, with its widespread projections to the PFC and NAc (Arias-Carrion & Poppel 2007; Haber 2003), will worsen task performance in healthy individuals as intrinsic dopamine levels in healthy individuals are likely to be optimised. Indeed, it has been shown that the relationship between dopamine and performance in many cognitive tasks,

especially those relying on working memory (Stuss & Knight 2002a), follows an inverted U-shaped function with both reducing or increasing the levels of dopamine leading to worsened task performance (Cools 2006; Cools & D'Esposito 2011; Robbins 2000; Williams & Goldman-Rakic 1995). In PD patients it has been demonstrated that while L-dopa administration might improve motor function it can move them away from their cognitive optimum (Gotham et al 1988; Rowe et al 2008), putatively by an 'overdosing' of the cortico-striatal loop involved in the task in question. Concordant with our findings, although within the confines of a very different task, L-dopa administration to Parkinson's patients impaired performance in probabilistic reversal learning and mainly in those which are signalled by negative feedback (Cools et al 2006; Cools et al 2001a). Furthermore, PD patients have been shown to have deficits in other tasks which require behavioural adaptation after rule changes (Cools et al 2001b; Gotham et al 1988). These behavioural deficits have been associated with abnormal activations in the cortex and striatum in PD patients, which were further dependant on whether positive or negative feedback were received (Monchi et al 2004), pointing to an influence of both dopaminergic status and feedback type on tasks of this ilk.

We additionally predicted a reciprocity in activation patterns between subcortical and cortical structures receiving major dopaminergic projections. We indeed observed such reciprocity with decreases in prefrontal activity concurrent with increases in caudate activity. This finding is in keeping with the known neurochemical reciprocity between the PFC and striatum whereby increases in dopamine in the PFC are associated with decreased dopamine in the basal ganglia

and vice versa (Pycock et al 1980; van Schouwenburg et al 2010), and also with the involvement of the caudate in tasks of this type (Clarke et al 2011). It may also be reflective of the different time courses of activity in the PFC and striatum which has been demonstrated in several tasks (Fujii & Graybiel 2005; Pasupathy & Miller 2005), and accords with theories suggesting a 'gating' role for the basal ganglia (Frank & Claus 2006; McNab & Klingberg 2008), and specifically a critical role for dopamine in this 'gating' (Miller & Cohen 2001).

It has been suggested that dopamine in the PFC and the basal ganglia regulate the balance between 2 functionally opponent processes, with the PFC dopamine regulating stability and the basal ganglia dopamine promoting cognitive flexibility (Cools 2008; van Schouwenburg et al 2010). Our data enriches these theories by suggesting that meta-stable representations of behavioural set are maintained in PFC under dopaminergic influence, and that the destruction and reestablishment of these cortical dynamics corresponding to a set switch, relies on a transient suppression of dopaminergically mediated activity following a hedonically surprising, negatively valenced outcome; a suppression of activity which we disrupt by exogenous L-dopa administration.

My central finding that dopamine administration abolishes the innate ability of healthy humans to alter behaviour in the face of negative outcomes, has great clinical relevance for understanding impulse control disorders observed in Parkinson's disease patients. These patients suffer from compulsive and impulsive behaviours as a result of dopamine replacement medication whereby they are unable to disengage with seemingly pointless activities such as compulsive

gardening and grooming even when these lead to negative life outcomes (Evans et al 2004; Evans et al 2009; McKeon et al 2007). A blunting or reversal of the normal suppression of activity could provide an explanation for this paradoxical maintenance of behavioural set in the face of negative outcomes. Here we propose possible neurobiological mechanisms for these behavioural disorders, and demonstrate a pervasive neurobiological role of dopamine in both stability and switching of responses which transcends action.

Chapter 7

Discussion

7.1 Summary of main results

Dopamine is a central neurotransmitter in the basal ganglia and influences many different aspects of behaviour. Among its many roles, dopamine has effects on both learning (Bayer & Glimcher 2005; Schultz 1998; Schultz et al 1997; Wise 2004; Wise & Rompre 1989) and on performance (i.e. on the expression of learnt behaviour) (Bardgett et al 2009; Berridge 2007; Boureau & Dayan 2011; Dickinson et al 2000; Mazzoni et al 2007; Niv 2007; Parkinson et al 2002; Salamone et al 2003). As described earlier in this thesis, these effects are confounded in many experiments involving dopamine manipulations (Cools et al 2007b; Frank et al 2007; Frank et al 2004; Pessiglione et al 2006).

In my first experiment, *Dopamine and performance in a reinforcement learning task – evidence from Parkinson's disease* (see Chapters 3 and 4), I was able to distinguish these factors by utilising the model afforded by Parkinson's Disease (PD) and performing a within-subject controlled drug manipulation. This allowed me to test whether the previously observed impairments in reinforcement learning in PD (Frank et al 2004; Knowlton et al 1996) were attributable to the effect of dopamine depletion on the learning process itself or on the expression of learning, i.e. on action performance.

By testing patients in different drug states I found that the main effect of dopamine was actually on the *performance* rather than the *learning* aspect of feedback-related learning, a finding which poses a challenge to many theories on the role of dopamine in learning (Bayer & Glimcher 2005; Frank et al 2007; Frank et al 2004;

O'Doherty et al 2003; Pessiglione et al 2006; Schultz et al 1997). At the neural level, the improved accuracy observed during the performance stage in the patients who were ON medication (irrelevant of which state they had been in when they learnt the contingencies) was associated with enhanced nucleus accumbens (NAc) and ventromedial prefrontal cortex (vmPFC) activity for the chosen cue value, an effect absent in the OFF medication state. The enhanced activity in the NAc and vmPFC in the ON medication state (which correlated with the improved behaviour) suggests that enhanced cue value representation underlies successful choice behaviour by improving patients' ability to choose the better outcome in novel contexts either by a more stable cue value representation or by improving the ability to compare the values of stimuli leading to such outcomes.

These findings are important for the clinical and neuroscientific community as they potentially clarify effects of dopamine and also allow an improved understanding of how the various roles of dopamine interact and overlap. This overlap has previously made it difficult to differentiate between dopamine's various roles, and I was able to achieve this differention by using the human model of dopamine depletion provided by PD.

My second experiment, *The effect of valence on movement: a study of bradykinesia in Parkinson's disease* (see Chapter 5) utilised the findings of the first experiment and went on to demonstrate a link between the cognitive and motor deficits observed in Parkinson's disease. A key result from my first experiment was that patients OFF medication are impaired at subsequently picking the most rewarding stimuli (thereby leading to lower accuracy in the task overall). In this first

experiment however, there were no specifically aversive outcomes (only outcomes which were more or less likely to be probabilistically correct). Consequently, I then questioned as to what effect dopamine depletion would have on specifically aversive outcomes and furthermore whether this would have an effect on bradykinesia, one of the main movement deficits found in PD.

In designing the second experiment, my aim was to establish whether the reward insensitivity in this group (i.e. PD patients OFF medication) would carry over into the motor domain by comparing the performance of PD subjects in rewarding scenarios to aversive ones. The impaired responsiveness to rewards which we observed in the first experiment when dopamine levels were low (due to disease), was not surprising given the pivotal role dopamine plays in reward (Bayer & Glimcher 2005; Day et al 2006; Schultz et al 1997). Indeed, the impaired adjustment of movement time in response to rewards has previously been shown in PD patients OFF medication (Moustafa et al 2008). What I aimed to demonstrate however, was that this deficit would be manifest in impaired motor speed to rewarding stimuli but less so to aversive stimuli.

The results confirmed our predictions in that I was able to clearly demonstrate a valence asymmetry in the movement time of PD patients, whereby there is a comparative failure to speed up in order to win rewards compared with an ability to speed up in order to avoid punishments. We also found that although the trial-by-trial response to rewards was impaired in the PD group, the trial-by-trial response to punishments was not similarly impacted, a finding which has not previously been demonstrated. Finally, we also found a detrimental effect of distractors which was

only evident for shock trials in PD patients, showing that sensitivity to distraction is also valence specific. These findings are important, as they demonstrate that movement time in PD is dependent on valence context, in other words whether movements are to harness a reward or avoid a punishment matters. I found that there is an asymmetry in patients' ability to move fast in response to rewards compared to punishments, demonstrating that even when tested in the context of everyday outcomes, bradykinesia is a variable, context-dependent deficit.

The above findings support the hypothesis that there is significant overlap between the motor and associative/limbic basal ganglia loops and provide empirical evidence that one can significantly impact on the other. It also supports hypotheses that the striatum and dopaminergic transmission are key to this influence (Haber 2003; Mogenson et al 1980). The results of my second experiment also have direct clinical relevance for PD patients. By showing that attention and implicit motivation are significant contributors to movement time in this group of patients, physical therapies in the future might be usefully adapted to take account of this fact leading to increased effectiveness and better outcomes.

In my third experiment, *Expectations and violations: Probing the role of dopamine in set shifting* (see Chapter 6), I turned my attention away from the effects of dopamine depletion and studied how boosting dopamine affects behaviour in the context of set shifting. For this experiment, using healthy human subjects, I isolated the effect of dopamine on set switching by specifically controlling for both its effects on motor vigour and for its effect on responses to violations of sensory predictions. I found that only subjects with normal dopaminergic function (i.e.

which have not been perturbed by administration of L-dopa), could effectively set switch in response to cues which carry negative valence. When these same subjects were given dopamine they lost this native ability, demonstrating that “overdosing” the dopaminergic system leads to deficits in responses to negative violations of expectations. Furthermore, when subjects were under the influence of L-dopa we observed a reversal of the normal pattern of behaviour whereby subjects became slightly (albeit not significantly) *better* at set switching in response to null outcomes compared to outcomes with negative valence, leading to a two way drug by valence interaction.

Effective set switching in the placebo group was associated with a larger decrease in PFC neuronal activity during unexpected losses (when the unexpected switch trial was contrasted with the fully expected pre-shift trial). In the L-dopa group this decrease in neuronal activity was absent on loss trials and paradoxically was only present in the null trials. These results support hypotheses that the relationship between dopamine and performance in many cognitive tasks, especially those relying on working memory (Stuss & Knight 2002), follow an inverted U-shaped function whereby the optimum level of performance exists at a certain level of dopaminergic stimulation. Any movement away from that peak (either by artificially reducing or increasing the levels of dopamine) leads to worsened task performance (Cools 2006; Cools & D'Esposito 2011; Gotham et al 1988; Robbins 2000; Rowe et al 2008; Williams & Goldman-Rakic 1995).

It is reasonable to assume that humans have evolved to have the optimal dopaminergic system for survival, with a fine balance being achieved between the

evolutionary need to avoid aversive cues, which could signal death or injury (with disastrous results which should therefore be avoided at even high costs), while maintaining appropriate responses to neutral or rewarding stimuli. In fact, the inability to learn correctly from punishments is a feature of many pathological lesions and conditions (Bechara 2005; Patterson & Newman 1993). The results of my research indicate that this fine balance can be altered by modulation of the dopaminergic system. It further demonstrates that: (a) low levels of dopamine such as are found in unmedicated PD patients is associated with poor responses, both in motor and cognitive tasks, to rewarding stimuli but a relatively preserved responses to aversive outcomes; (b) augmented dopamine states, such as when PD patients are given dopamine replacement therapy (DRT) or when healthy subjects are given L-dopa, are associated with improved performance in relation to null or positive outcomes but can lead to less effective responses to loss outcomes; and (c) in healthy controls who have normal, unperturbed dopaminergic systems, there is a bias towards better responding to negative outcomes. Evidence for this comes from the second experiment (Chapter 5) in which the movement times in the control group were faster when subjects avoided aversive outcomes compared to when they tried to reap rewarding outcomes which demonstrates that the avoidance (moving away from) pain is of greater importance than a successful pursuit of reward. Further evidence for this evolutionary asymmetry comes from the placebo group in the third experiment (Chapter 6) whereby this group showed more effective set shifting in response to losses compared with null outcomes, when all other major factors were controlled for. It appears therefore that low dopamine states exacerbate this evolutionarily proscribed asymmetry, where there is a bias

towards better loss avoidance, an asymmetry eliminated under artificially raised dopamine states.

Note that all of the effects I have found are separate from those on learning. In the first experiment, when we attempted to separate learning from performance, we found that dopamine exerted its effects primarily on the expression of learning rather than on learning itself. In the second experiment no significant learning took place. In the third experiment we examined set shifting as opposed to directly testing learning.

7.2 Implications for theories of dopamine function

As mentioned above, dopamine has been implicated in many different aspects of cogntion including learning (Montague et al 2004; Schultz et al 1997), motivation and vigour (Salamone & Correa 2002; Ungerstedt 1971) and action selection and movement (O'Doherty 2004; Samejima et al 2005). In the experiments in this thesis I sought to disambiguate several of these functions from one another in order to attempt to improve the understanding of the function of dopamine. The results of the first experiment indicate a central role for dopamine in the performance aspect of reinforcement learning rather then on learning. Our results indicate that levels of dopamine do not impact on actual learning but rather on the expression of that learning in contrary to many previous accounts (Bayer & Glimcher 2005; Frank et al 2007; Frank et al 2004). One of the reasons for the differences between our results and previous findings may be in part due to many previous experiments confounding learning and performance. There are however many different types of learning which lead to various behavioural outputs (Dickinson et al 2000; Palmiter 2008) and as such the lack of an effect of dopamine on learning in our specific experiment does not of course rule out a role for dopamine entirely in learning. I do believe however that this experiment should act as an example for future work of the importance of using tasks optimised for isolating specific functions to avoid the possibility of confounding the various roles of dopamine. In relation to the seminal work of Schultz (Schultz et al 1997) and the reward prediction error findings in animals, our data cannot specifically address the issue of whether prediction error related activity is the main driving force behind learning or instead occurs as a

consequence of learning as proposed by Berridge and Robinson (Berridge 2007; Berridge & Robinson 1998) only that we could not isolate a direct effect of manipulating dopamine on learning.

Dopamine has been widely implicated directly in reward processing (Wise 1978) in both animals and in humans. This is an extra complicating factor when attempting to disambiguate the roles of dopamine from each other. As a consequence of this fact, the second experiment sought to detect whether the effect of valence would carry influence motor output directly, while the third experiment directly examined the role of valence in the ability of humans to set shift in both normal and boosted dopamine states. We found that valence is in fact a critical factor in the function of dopamine and even when other factors are controlled for it exerts one of its main effect in this domain. This in turn has both implications for PD (see next section) and for understanding human motivated behaviour in general.

7.3 Implications for PD

The research I have carried out and have detailed in this thesis makes a potentially important contribution to our understanding of the origins of some of the deficits observed in Parkinson's disease. The results of the first experiment demonstrate that low dopamine levels detrimentally impact value representation in decision making in novel contexts, an ability restored by dopamine treatment. The key clinical implication of this finding is that patients' decision making abilities are affected by dopamine replacement therapy.

Most patients and physicians are unaware that administration of DRT can have an effect on even the simplest of implicit choices and that this effect could have detrimental consequences on both the major and minor everyday decisions which patients make. In addition, the improvement in reward related performance observed in patients who are on their DRT provides a potential explanation of some of the compulsive behaviours observed in patients who overuse dopamine medication. From this work, we can infer that even in PD patients who do not suffer with overt impulse control disorders, dopaminergic medication is having an effect on the choices they make and on their ability to respond correctly to stimuli with different valence characteristics. This has important implications to PD and its treatment and has relevance to the day-to-day choices patients make. Patients and the physicians who treat them need to be more aware of this effect of medication on cognition.

In the second experiment, I link the different responses to positive and negative outcomes to the movement deficit observed in PD patients and in doing so demonstrate that the outcome of a movement (i.e. reap a positive outcome or avoid a negative outcome) impacts the speed of movement and that this asymmetry is more marked in PD patients than in controls. This demonstrates that dopamine depletion has a lesser impact on responses to punishments compared to rewards. This finding provides a potential explanation for the paradoxical kinesis observed in PD patients, whereby patients are suddenly able to move at near normal speeds, usually in extreme aversive contexts, and also demonstrated that everyday outcomes can have very different effects on movement time in PD. The understanding that Parkinson's disease patients' ability to move might depend on factors that may have previously not been considered as relevant could lead to more effective future strategies in physical therapy.

Finally, in the third experiment, I have shown that overdosing the dopaminergic system in healthy individuals leads to a performance decrement in set shifting as well as a loss of the innate ability to avoid aversive outcomes. Overdosing the dopaminergic system actually leads to subjects responding abnormally to null outcomes. This is in keeping with theories on the inverted U-shaped function whereby the optimum level of performance exists at a certain level of dopaminergic stimulation, and any movement away from that peak (either by artificially reducing or increasing the levels of dopamine) leads to worsened task performance. This has been clearly demonstrated in several studies on PD patients (Cools et al 2001a; Cools & D'Esposito 2011) and has been proposed to be due to the overdosing of the

ventral striatal orbitofrontal circuity which is relatively unaffected in early PD (Cools et al 2001a). Here we propose a central role for valence in behaviour, and that a similar U-shaped function may exist in this domain as well and that depletion of dopamine (such as is found in the PD state) enhances the effect of negative outcomes on behaviour and boosting dopamine, either in healthy or parkinsonian patients enhances the effect of cues which yield positive outcomes.

A blunting or reversal of the normal response to outcomes which carry negative valence could provide an explanation for the paradoxical maintenance of behavioural sets in the face of negative outcomes observed in PD patients who compulsively use dopaminergic medication (Evans et al 2004) and potentially in addictive disorders in general (Bechara 2005; Patterson & Newman 1993). Impulse control disorders in PD are common with prevalence rates of 13.6%. They consist of behaviours such as pathological gambling, compulsive buying, compulsive sexual behavior, and binge or compulsive eating and have a strong association with the use of dopaminergic medication. Dopamine agonist treatment in PD is associated with a 2-3.5 fold increase in these behaviours (Voon et al 2011; Weintraub et al 2010). It is clear that they are important disorders which are influenced by dopaminergic medication. Here we propose a possible mechanism for these disorders.

7.4 Conclusions and future directions

The aim of my research which I have detailed in this thesis was to probe the effects of dopamine on cognition and to attempt, with this knowledge, to understand more about the cognitive effects of Parkinson's disease, a disorder characterised by loss of dopamine. To this end I utilised the model afforded by these patients and tested them in different medication states with the hope of elucidating the effects of dopamine in the various cognitive domains. As detailed previously however, using Parkinson's disease as a model for dopamine depletion is problematic in several ways. Firstly other neurotransmitters such as serotonin and acetylcholine are affected in PD as well as the fact that PD patients tend to be older, a factor also associated with decline in other neurotransmitters. In addition, Parkinson's disease patients are a very hetrogenous group, thereby making inferences about this population difficult to apply across the board. Despite all these caveats, Parkinson's disease remains the only viable model of dopamine depletion in humans and I therefore believe that much about both the disease and about the function of dopamine in the healthy brain can be understood by testing these patients.

There are several key findings that arise from the research which I have undertaken. By disambiguating learning from performance in a key reinforcement learning task I have been able to demonstrate that dopamine exerts a critical effect on performance which is separate and more important than the effect it has on learning. This finding provides a challenge to much of the learning literature and provides important insights into the function of dopamine.

A further central finding from this work is that dopamine appears to exert an effect on responsiveness to cues with different valence contexts, with boosted dopamine states increasing responsiveness to positive or null outcomes, while unperturbed or depleted dopamine levels associated with better responsiveness to outcomes with negative valence. This finding was evident even when the effects of action were controlled for. This has implications in the search for a better understanding of compulsive disorders in PD, which are associated with the use of dopaminergic medication (Weintraub et al 2010).

Another key finding from my research is to provide evidence to support the hypotheses that activity in one loop of the basal ganglia may be influenced by activity in a separate loop (Haber 2003; Haber & Knutson 2010; Redgrave et al 2010). In the second experiment I demonstrate, by comparing Parkinson's disease patients and controls, that the cognitive loops processing reward and punishment exert a distinct impact on motor performance. Furthermore, I show that bradykinetic movements in PD are dependent on context, and that in the dopamine depleted state, responses to rewards are impaired while responses to punishments are relatively well maintained– which is both a novel and important finding. In the third experiment, I provide empirical evidence for the reciprocity of response patterns between the mesocortical projections to the vmPFC, and the nigrostriatal projections to the caudate and putamen. These findings, put together, persuasively demonstrate that these loops are both influenced and modulated by one another, and support the hypotheses suggesting that limbic inputs influence motor output via the striatal loops and are under the influence of the dopaminergic system (Haber 2003; Mogenson et al 1980; Roitman et al 2005).

In sum, I hope that the findings from this work will significantly contribute to the growing knowledge base surrounding the function of dopamine. These findings could have important implications:

(a) for PD patients by providing a better understanding of the cognitive deficits manifest in this disease thereby leading to better management strategies in the future, as well as a greater understanding of the implications of the use of certain forms of current treatment;

(b) as a basis for future work exploring and building a better understanding of the functions of dopamine in the healthy brain; and

(c) by pointing to possible neurobiological mechanisms which underlie compulsive disorders.

In order to build on this research, in the future I would be interested in focussing on the modulatory role of dopamine in different types of learning by by using careful experimental designs.

Given the influence of almost all neurotransmitters on action as well as cognition, widespread utilisation of the Go/NoGo task design will allow for better and more accurate control of the effects of action, eliminating as much as possible this very important confound.

I would also be interested in examining further the effect of dopamine replacement in PD on responses to punishments and rewards by examining motor performance in a similar fashion to the experiment in chapter 5. Furthermore , it would be of interest to examine the interaction between medication status and graded rewards

and punishments in a large sample size of PD patients and matched controls to further clarify these effects.

Another direction I would be keen to pursue is to examine the performance of PD patients on my final experiment, in which I have so far only tested healthy volunteers on placebo and on L-dopa. This would allow me to examine the influence of valence on set shifting in PD while fully controlling for motor output by using this Go/NoGo design. In fact, results from an experiment such as this and those detailed above would provide important support for some of the conclusions of this thesis that have led to the proposal that above many other factors, dopamine has a central influence on the modulation of cognitive performance in relation to different valence outcomes.

Publications arising from this work

1. Shiner T, Seymour B, Wunderlich W, Hill C, Bhatia KP, Dayan P, Dolan RJ. DOPAMINE AND PERFORMANCE IN A REINFORCEMENT LEARNING TASK-EVIDENCE FROM PARKINSON'S DISEASE. *Accepted to Brain*.
2. Shiner T, Seymour B, Symmonds M, Dayan P, Bhatia KP, Dolan RJ. THE EFFECT OF VALENCE ON MOVEMENT: A STUDY OF BRADYKINESIA IN PARKINSON'S DISEASE. *In prep*.
3. Shiner T, Symmonds M, Guitart-Masip M, Fleming S, Friston K, Dolan RJ. EXPECTATIONS AND VIOLATIONS: PROBING THE ROLE OF DOPAMINE IN SET SHIFTING. *In prep*.
4. Sharot T, Shiner T, Brown AC, Fan J. DOPAMINE ENHANCES EXPECTATION OF PLEASURE IN HUMANS. *Current Biology. 2009 Dec 29;19(24):2077-80.*
5. Sharot T, Shiner T, Dolan RJ. HOW EXPERIENCE & CHOICE SHAPE EXPECTED AVERSIVE OUTCOMES. *Journal of Neuroscience. 2010 Jul 7;30(27):9209-15.*
6. Pine A, Shiner T, Seymour B, Curran HV, Dolan RJ. DOPAMINE, TIME AND IMPULSIVITY IN HUMANS. *Journal of Neuroscience. 2010 Jun 30;30(26):8888-96.*

References

Aarsland D, Bronnick K, Larsen JP, Tysnes OB, Alves G. 2009. Cognitive impairment in incident, untreated Parkinson disease: the Norwegian ParkWest study. *Neurology* 72:1121-6

Aizman O, Brismar H, Uhlen P, Zettergren E, Levey AI, et al. 2000. Anatomical and physiological evidence for D1 and D2 dopamine receptor colocalization in neostriatal neurons. *Nat Neurosci* 3:226-30

Albers RW, Brady ST, Price DL, eds. 2006. *Basic Neurochemistry. Molecular, cellular and medical aspects*: Elsevier

Albin RL, Young AB, Penney JB. 1989. The functional anatomy of basal ganglia disorders. *Trends Neurosci* 12:366-75

Alexander GE, DeLong MR, Strick PL. 1986. Parallel organization of functionally segregated circuits linking basal ganglia and cortex. *Annu Rev Neurosci* 9:357-81

Antonini A, Leenders KL, Reist H, Thomann R, Beer HF, Locher J. 1993. Effect of age on D2 dopamine receptors in normal human brain measured by positron emission tomography and 11C-raclopride. *Arch Neurol* 50:474-80

Arias-Carrion O, Poppel E. 2007. Dopamine, learning, and reward-seeking behavior. *Acta Neurobiol Exp (Wars)* 67:481-8

Asaad WF, Rainer G, Miller EK. 2000. Task-specific neural activity in the primate prefrontal cortex. *J Neurophysiol* 84:451-9

Ashburner J, Friston KJ. 2005. Unified segmentation. *Neuroimage* 26:839-51

Baik JH, Picetti R, Saiardi A, Thiriet G, Dierich A, et al. 1995. Parkinsonian-like locomotor impairment in mice lacking dopamine D2 receptors. *Nature* 377:424-8

Bardgett ME, Depenbrock M, Downs N, Points M, Green L. 2009. Dopamine modulates effort-based decision making in rats. *Behav Neurosci* 123:242-51

Bayer HM, Glimcher PW. 2005. Midbrain dopamine neurons encode a quantitative reward prediction error signal. *Neuron* 47:129-41

Bechara A. 2005. Decision making, impulse control and loss of willpower to resist drugs: a neurocognitive perspective. *Nat Neurosci* 8:1458-63

Beck AT, Ward CH, Mendelson M, Mock J, Erbaugh J. 1961. An inventory for measuring depression. *Arch Gen Psychiatry* 4:561-71

Becker JB, Cha JH. 1989. Estrous cycle-dependent variation in amphetamine-induced behaviors and striatal dopamine release assessed with microdialysis. *Behav Brain Res* 35:117-25

Beninger RJ. 1983. The role of dopamine in locomotor activity and learning. *Brain Res* 287:173-96

Berardelli A, Dick JP, Rothwell JC, Day BL, Marsden CD. 1986. Scaling of the size of the first agonist EMG burst during rapid wrist movements in patients with Parkinson's disease. *J Neurol Neurosurg Psychiatry* 49:1273-9

Berke JD, Hyman SE. 2000. Addiction, dopamine, and the molecular mechanisms of memory. *Neuron* 25:515-32

Berridge KC. 2007. The debate over dopamine's role in reward: the case for incentive salience. *Psychopharmacology (Berl)* 191:391-431

Berridge KC. 2009. 'Liking' and 'wanting' food rewards: brain substrates and roles in eating disorders. *Physiol Behav* 97:537-50

Berridge KC, Robinson TE. 1998. What is the role of dopamine in reward: hedonic impact, reward learning, or incentive salience? *Brain Res Brain Res Rev* 28:309-69

Blin O, Ferrandez AM, Serratrice G. 1990. Quantitative analysis of gait in Parkinson patients: increased variability of stride length. *J Neurol Sci* 98:91-7

Bloom FE, Kupfer DJ, eds. 1995. *Psychopharmacology:The Fourth Generation of Progress*. New York: Raven Press

Bodi N, Keri S, Nagy H, Moustafa A, Myers CE, et al. 2009. Reward-learning and the novelty-seeking personality: a between- and within-subjects study of the effects of dopamine agonists on young Parkinson's patients. *Brain* 132:2385-95

Boorman ED, Behrens TE, Woolrich MW, Rushworth MF. 2009. How green is the grass on the other side? Frontopolar cortex and the evidence in favor of alternative courses of action. *Neuron* 62:733-43

Boureau YL, Dayan P. 2011. Opponency revisited: competition and cooperation between dopamine and serotonin. *Neuropsychopharmacology* 36:74-97

Braak H, Ghebremedhin E, Rub U, Bratzke H, Del Tredici K. 2004. Stages in the development of Parkinson's disease-related pathology. *Cell Tissue Res* 318:121-34

Brown RG, Marsden CD. 1988. Internal versus external cues and the control of attention in Parkinson's disease. *Brain* 111 (Pt 2):323-45

Buxton RB. 2009. *Introduction to functional magnetic resonance imaging*. New Yorki: Cambridge university press

Cannon CM, Palmiter RD. 2003. Reward without dopamine. *J Neurosci* 23:10827-31

Carr DB, Sesack SR. 2000. Projections from the rat prefrontal cortex to the ventral tegmental area: target specificity in the synaptic associations with mesoaccumbens and mesocortical neurons. *J Neurosci* 20:3864-73

Centonze D, Picconi B, Gubellini P, Bernardi G, Calabresi P. 2001. Dopaminergic control of synaptic plasticity in the dorsal striatum. *Eur J Neurosci* 13:1071-7

Clarke HF, Hill GJ, Robbins TW, Roberts AC. 2011. Dopamine, but not serotonin, regulates reversal learning in the marmoset caudate nucleus. *J Neurosci* 31:4290-7

Cools R. 2006. Dopaminergic modulation of cognitive function-implications for L-DOPA treatment in Parkinson's disease. *Neurosci Biobehav Rev* 30:1-23

Cools R. 2008. Role of dopamine in the motivational and cognitive control of behavior. *Neuroscientist* 14:381-95

Cools R, Altamirano L, D'Esposito M. 2006. Reversal learning in Parkinson's disease depends on medication status and outcome valence. *Neuropsychologia* 44:1663-73

Cools R, Barker RA, Sahakian BJ, Robbins TW. 2001a. Enhanced or impaired cognitive function in Parkinson's disease as a function of dopaminergic medication and task demands. *Cereb Cortex* 11:1136-43

Cools R, Barker RA, Sahakian BJ, Robbins TW. 2001b. Mechanisms of cognitive set flexibility in Parkinson's disease. *Brain* 124:2503-12

Cools R, D'Esposito M. 2011. Inverted-u-shaped dopamine actions on human working memory and cognitive control. *Biol Psychiatry* 69:e113-25

Cools R, Lewis SJ, Clark L, Barker RA, Robbins TW. 2007a. L-DOPA disrupts activity in the nucleus accumbens during reversal learning in Parkinson's disease. *Neuropsychopharmacology* 32:180-9

Cools R, Miyakawa A, Sheridan M, D'Esposito M. 2010. Enhanced frontal function in Parkinson's disease. *Brain* 133:225-33

Cools R, Nakamura K, Daw ND. 2011. Serotonin and dopamine: unifying affective, activational, and decision functions. *Neuropsychopharmacology* 36:98-113

Cools R, Roberts AC, Robbins TW. 2008a. Serotoninergic regulation of emotional and behavioural control processes. *Trends Cogn Sci* 12:31-40

Cools R, Robinson OJ, Sahakian B. 2008b. Acute tryptophan depletion in healthy volunteers enhances punishment prediction but does not affect reward prediction. *Neuropsychopharmacology* 33:2291-9

Cools R, Sheridan M, Jacobs E, D'Esposito M. 2007b. Impulsive personality predicts dopamine-dependent changes in frontostriatal activity during component processes of working memory. *J Neurosci* 27:5506-14

Creese I, Sibley DR, Hamblin MW, Leff SE. 1983. The classification of dopamine receptors: relationship to radioligand binding. *Annu Rev Neurosci* 6:43-71

Critchley M. 1929. Arteriosclerotic Parkinsonism. *Brain* 52:23-83

Crofts HS, Dalley JW, Collins P, Van Denderen JC, Everitt BJ, et al. 2001. Differential effects of 6-OHDA lesions of the frontal cortex and caudate nucleus on the ability to acquire an attentional set. *Cereb Cortex* 11:1015-26

Czernecki V, Pillon B, Houeto JL, Pochon JB, Levy R, Dubois B. 2002. Motivation, reward, and Parkinson's disease: influence of dopatherapy. *Neuropsychologia* 40:2257-67

Daw ND, Kakade S, Dayan P. 2002. Opponent interactions between serotonin and dopamine. *Neural Netw* 15:603-16

Daw ND, Niv Y, Dayan P. 2005. Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nat Neurosci* 8:1704-11

Day JJ, Wheeler RA, Roitman MF, Carelli RM. 2006. Nucleus accumbens neurons encode Pavlovian approach behaviors: evidence from an autoshaping paradigm. *Eur J Neurosci* 23:1341-51

Dayan P, Niv Y, Seymour B, Daw ND. 2006. The misbehavior of value and the discipline of the will. *Neural Netw* 19:1153-60

de Rijk MC, Launer LJ, Berger K, Breteler MM, Dartigues JF, et al. 2000. Prevalence of Parkinson's disease in Europe: A collaborative study of population-based cohorts. Neurologic Diseases in the Elderly Research Group. *Neurology* 54:S21-3

de Wit S, Barker RA, Dickinson AD, Cools R. 2011. Habitual versus Goal-directed Action Control in Parkinson Disease. *J Cogn Neurosci* 23:1218-29

Deakin JF, Graeff FG. 1991. 5-HT and mechanisms of defence. *J Psychopharmacol* 5:305-15

Dickinson A, Smith J, Mirenowicz J. 2000. Dissociation of Pavlovian and instrumental incentive learning under dopamine antagonists. *Behav Neurosci* 114:468-83

Dodd ML, Klos KJ, Bower JH, Geda YE, Josephs KA, Ahlskog JE. 2005. Pathological gambling caused by drugs used to treat Parkinson disease. *Arch Neurol* 62:1377-81

Draganski B, Bhatia KP. 2010. Brain structure in movement disorders: a neuroimaging perspective. *Curr Opin Neurol* 23:413-9

Draganski B, Kherif F, Kloppel S, Cook PA, Alexander DC, et al. 2008. Evidence for segregated and integrative connectivity patterns in the human Basal Ganglia. *J Neurosci* 28:7143-52

Drago J, Gerfen CR, Lachowicz JE, Steiner H, Hollon TR, et al. 1994. Altered striatal function in a mutant mouse lacking D1A dopamine receptors. *Proc Natl Acad Sci U S A* 91:12564-8

Dreher JC, Schmidt PJ, Kohn P, Furman D, Rubinow D, Berman KF. 2007. Menstrual cycle phase modulates reward-related neural function in women. *Proc Natl Acad Sci U S A* 104:2465-70

Dubois B, Pillon B. 1997. Cognitive deficits in Parkinson's disease. *J Neurol* 244:2-8

Durstewitz D, Seamans JK, Sejnowski TJ. 2000. Neurocomputational models of working memory. *Nat Neurosci* 3 Suppl:1184-91

Ebadi M, Pfeiffer RF, eds. 2005. *Parkinson's Disease*: CRC press

Edwards MQ, N. Bhatia, K. 2008. Parkinson's disease and other movement disorders. Oxford University Press

Eppinger B, Hammerer D, Li SC. 2011. Neuromodulation of reward-based learning and decision making in human aging. *Ann N Y Acad Sci* 1235:1-17

Erixon-Lindroth N, Farde L, Wahlin TB, Sovago J, Halldin C, Backman L. 2005. The role of the striatal dopamine transporter in cognitive aging. *Psychiatry Res* 138:1-12

Evans AH, Katzenschlager R, Paviour D, O'Sullivan JD, Appel S, et al. 2004. Punding in Parkinson's disease: its relation to the dopamine dysregulation syndrome. *Mov Disord* 19:397-405

Evans AH, Strafella AP, Weintraub D, Stacy M. 2009. Impulsive and compulsive behaviors in Parkinson's disease. *Mov Disord* 24:1561-70

Factor SA, Weiner WJ, eds. 2008. *Parkinson's disease. Diagnosis and clinical management*: Demos medical publishing

Fahn S ER, Committee MotUD, ed. 1987. *Unified Parkinson's disease rating scale*: McMillan Health Care Information

Faro S, Mohamed F, eds. 2010. *BOLD fMRI: A Guide to Functional Imaging for Neuroscientists*. New York: Springer

Fitzgerald TH, Seymour B, Bach DR, Dolan RJ. 2010. Differentiable neural substrates for learned and described value and risk. *Curr Biol* 20:1823-9

FitzGerald TH, Seymour B, Dolan RJ. 2009. The role of human orbitofrontal cortex in value comparison for incommensurable objects. *J Neurosci* 29:8388-95

Folstein MF, Folstein SE, McHugh PR. 1975. "Mini-mental state". A practical method for grading the cognitive state of patients for the clinician. *J Psychiatr Res* 12:189-98

Foltynie T, Brayne CE, Robbins TW, Barker RA. 2004. The cognitive ability of an incident cohort of Parkinson's patients in the UK. The CamPaIGN study. *Brain* 127:550-60

Frackowiak RSJ, Friston KJ, Frith C, Dolan RJ, Price CJ, et al, eds. 2003. *Human Brain Function*: Academic press

Frank MJ, Claus ED. 2006. Anatomy of a decision: striato-orbitofrontal interactions in reinforcement learning, decision making, and reversal. *Psychol Rev* 113:300-26

Frank MJ, Samanta J, Moustafa AA, Sherman SJ. 2007. Hold your horses: impulsivity, deep brain stimulation, and medication in parkinsonism. *Science* 318:1309-12

Frank MJ, Seeberger LC, O'Reilly R C. 2004. By carrot or by stick: cognitive reinforcement learning in parkinsonism. *Science* 306:1940-3

Friston K, Ashburner J, Kiebel S, Nichols T, Penny W, eds. 2006. *Statistical Parametric Mapping: The Analysis of Functional Brain Images*: Academic press

Friston KJ. 1997. Transients, metastability, and neuronal dynamics. *Neuroimage* 5:164-71

Friston KJ, Dolan RJ. 2010. Computational and dynamic models in neuroimaging. *Neuroimage* 52:752-65

Fujii N, Graybiel AM. 2005. Time-varying covariance of neural activities recorded in striatum and frontal cortex as monkeys perform sequential-saccade tasks. *Proc Natl Acad Sci U S A* 102:9032-7

Funahashi S, Bruce CJ, Goldman-Rakic PS. 1989. Mnemonic coding of visual space in the monkey's dorsolateral prefrontal cortex. *J Neurophysiol* 61:331-49

Fuster JM. 2001. The prefrontal cortex--an update: time is of the essence. *Neuron* 30:319-33

Fuster JM, Alexander GE. 1971. Neuron activity related to short-term memory. *Science* 173:652-4

Gallin J, Ognibene F, eds. 2007. *Principals and practice of clinical research*: Academic press

Geisler S, Zahm DS. 2005. Afferents of the ventral tegmental area in the rat-anatomical substratum for integrative functions. *J Comp Neurol* 490:270-94

Gotham AM, Brown RG, Marsden CD. 1988. 'Frontal' cognitive function in patients with Parkinson's disease 'on' and 'off' levodopa. *Brain* 111 (Pt 2):299-321

Gottfried JA, O'Doherty J, Dolan RJ. 2003. Encoding predictive reward value in human amygdala and orbitofrontal cortex. *Science* 301:1104-7

Grace AA. 1991. Phasic versus tonic dopamine release and the modulation of dopamine system responsivity: a hypothesis for the etiology of schizophrenia. *Neuroscience* 41:1-24

Grace AA, Floresco SB, Goto Y, Lodge DJ. 2007. Regulation of firing of dopaminergic neurons and control of goal-directed behaviors. *Trends Neurosci* 30:220-7

Graef S, Biele G, Krugel LK, Marzinzik F, Wahl M, et al. 2010. Differential influence of levodopa on reward-based learning in Parkinson's disease. *Front Hum Neurosci* 4:169

Haber SN. 2003. The primate basal ganglia: parallel and integrative networks. *J Chem Neuroanat* 26:317-30

Haber SN, Knutson B. 2010. The reward circuit: linking primate anatomy and human imaging. *Neuropsychopharmacology* 35:4-26

Hallett M, Khoshbin S. 1980. A physiological mechanism of bradykinesia. *Brain* 103:301-14

Harrison AA, Everitt BJ, Robbins TW. 1999. Central serotonin depletion impairs both the acquisition and performance of a symmetrically reinforced go/no-go conditional visual discrimination. *Behav Brain Res* 100:99-112

Hausdorff JM, Balash J, Giladi N. 2003. Effects of cognitive challenge on gait variability in patients with Parkinson's disease. *J Geriatr Psychiatry Neurol* 16:53-8

Heller R, Stanley D, Yekutieli D, Rubin N, Benjamini Y. 2006. Cluster-based analysis of FMRI data. *Neuroimage* 33:599-608

Henson R. 2008. http://imaging.mrc-cbu.cam.ac.uk/imaging/DesignEfficiency

Hikida T, Kimura K, Wada N, Funabiki K, Nakanishi S. 2010. Distinct roles of synaptic transmission in direct and indirect striatal pathways to reward and aversive behavior. *Neuron* 66:896-907

Hnasko TS, Sotak BN, Palmiter RD. 2005. Morphine reward in dopamine-deficient mice. *Nature* 438:854-7

Hoehn MM, Yahr MD. 1967. Parkinsonism: onset, progression and mortality. *Neurology* 17:427-42

Hollander E. 1998. Treatment of obsessive-compulsive spectrum disorders with SSRIs. *Br J Psychiatry Suppl*:7-12

Hollerman JR, Schultz W. 1998. Dopamine neurons report an error in the temporal prediction of reward during learning. *Nat Neurosci* 1:304-9

Hornykiewicz O. 1998. Biochemical aspects of Parkinson's disease. *Neurology* 51:S2-9

Huettel SA, Song AW, McCarthy G. 2009. *Functional Magnetic Resonance Imaging* Sinauer Associates

Ikemoto S, Panksepp J. 1999. The role of nucleus accumbens dopamine in motivated behavior: a unifying interpretation with special reference to reward-seeking. *Brain Res Brain Res Rev* 31:6-41

Jezzard P, Matthews PM, Smith SM, eds. 2003. *Functional MRI: An Introduction to Methods*: Oxford University Press

Jhou TC, Fields HL, Baxter MG, Saper CB, Holland PC. 2009. The rostromedial tegmental nucleus (RMTg), a GABAergic afferent to midbrain dopamine neurons, encodes aversive stimuli and inhibits motor responses. *Neuron* 61:786-800

Jocham G, Klein TA, Ullsperger M. 2011. Dopamine-mediated reinforcement learning signals in the striatum and ventromedial prefrontal cortex underlie value-based choices. *J Neurosci* 31:1606-13

Joel D, Niv Y, Ruppin E. 2002. Actor-critic models of the basal ganglia: new anatomical and computational perspectives. *Neural Netw* 15:535-47

Josephs O, Henson RN. 1999. Event-related functional magnetic resonance imaging: modelling, inference and optimization. *Philos Trans R Soc Lond B Biol Sci* 354:1215-28

Kable JW, Glimcher PW. 2009. The neurobiology of decision: consensus and controversy. *Neuron* 63:733-45

Khor SP, Hsu A. 2007. The pharmacokinetics and pharmacodynamics of levodopa in the treatment of Parkinson's disease. *Curr Clin Pharmacol* 2:234-43

Kimura M, Aosaki T, Hu Y, Ishida A, Watanabe K. 1992. Activity of primate putamen neurons is selective to the mode of voluntary movement: visually guided, self-initiated or memory-guided. *Exp Brain Res* 89:473-7

Kincaid AE, Zheng T, Wilson CJ. 1998. Connectivity and convergence of single corticostriatal axons. *J Neurosci* 18:4722-31

Knowlton BJ, Mangels JA, Squire LR. 1996. A neostriatal habit learning system in humans. *Science* 273:1399-402

Koller W. 2007. *Parkinson's disease and related disorders*

Koller W, Melamed E. 2007a. *Parkinson's disease and related disorders part 1*: Elsevier

Koller WC, Melamed E, eds. 2007b. *Parkinson's disease and related disorders - part 2*, Vols. 84: Elsevier

Koller WC, Rueda MG. 1998. Mechanism of action of dopaminergic agents in Parkinson's disease. *Neurology* 50:S11-4; discussion S44-8

Kulisevsky J, Avila A, Barbanoj M, Antonijoan R, Berthier ML, Gironell A. 1996. Acute effects of levodopa on neuropsychological performance in stable and fluctuating Parkinson's disease patients at different levodopa plasma levels. *Brain* 119 (Pt 6):2121-32

Lees AJ, Smith E. 1983. Cognitive deficits in the early stages of Parkinson's disease. *Brain* 106 (Pt 2):257-70

Levant B. 1997. The D3 dopamine receptor: neurobiology and potential clinical relevance. *Pharmacol Rev* 49:231-52

Lewis GN, Byblow WD, Walt SE. 2000. Stride length regulation in Parkinson's disease: the use of extrinsic, visual cues. *Brain* 123 (Pt 10):2077-90

Lex B, Hauber W. 2010. The role of nucleus accumbens dopamine in outcome encoding in instrumental and Pavlovian conditioning. *Neurobiol Learn Mem* 93:283-90

Logothetis NK. 2003. The underpinnings of the BOLD functional magnetic resonance imaging signal. *J Neurosci* 23:3963-71

Logothetis NK, Pauls J, Augath M, Trinath T, Oeltermann A. 2001. Neurophysiological investigation of the basis of the fMRI signal. *Nature* 412:150-7

Luk CH, Wallis JD. 2009. Dynamic encoding of responses and outcomes by neurons in medial prefrontal cortex. *J Neurosci* 29:7526-39

Maia TV, Frank MJ. 2011. From reinforcement learning models to psychiatric and neurological disorders. *Nat Neurosci* 14:154-62

Matamales M, Bertran-Gonzalez J, Salomon L, Degos B, Deniau JM, et al. 2009. Striatal medium-sized spiny neurons: identification by nuclear staining and study of neuronal subpopulations in BAC transgenic mice. *PLoS One* 4:e4770

Matsumoto K, Suzuki W, Tanaka K. 2003. Neuronal correlates of goal-based motor selection in the prefrontal cortex. *Science* 301:229-32

Mazziotta JC, Toga AW, Evans A, Fox P, Lancaster J. 1995. A probabilistic atlas of the human brain: theory and rationale for its development. The International Consortium for Brain Mapping (ICBM). *Neuroimage* 2:89-101

Mazzoni P, Hristova A, Krakauer JW. 2007. Why don't we move faster? Parkinson's disease, movement vigor, and implicit motivation. *J Neurosci* 27:7105-16

McClure SM, Berns GS, Montague PR. 2003a. Temporal prediction errors in a passive learning task activate human striatum. *Neuron* 38:339-46

McClure SM, Daw ND, Montague PR. 2003b. A computational substrate for incentive salience. *Trends Neurosci* 26:423-8

McFarland NR, Haber SN. 2002. Thalamic relay nuclei of the basal ganglia form both reciprocal and nonreciprocal cortical connections, linking multiple frontal cortical areas. *J Neurosci* 22:8117-32

McKeon A, Josephs KA, Klos KJ, Hecksel K, Bower JH, et al. 2007. Unusual compulsive behaviors primarily related to dopamine agonist therapy in Parkinson's disease and multiple system atrophy. *Parkinsonism Relat Disord* 13:516-9

McNab F, Klingberg T. 2008. Prefrontal cortex and basal ganglia control access to working memory. *Nat Neurosci* 11:103-7

Mehta MA, Manes FF, Magnolfi G, Sahakian BJ, Robbins TW. 2004. Impaired set-shifting and dissociable effects on tests of spatial working memory following the dopamine D2 receptor antagonist sulpiride in human volunteers. *Psychopharmacology (Berl)* 176:331-42

Miller EK, Cohen JD. 2001. An integrative theory of prefrontal cortex function. *Annu Rev Neurosci* 24:167-202

Missale C, Nash SR, Robinson SW, Jaber M, Caron MG. 1998. Dopamine receptors: from structure to function. *Physiol Rev* 78:189-225

Mogenson GJ, Jones DL, Yim CY. 1980. From motivation to action: functional interface between the limbic system and the motor system. *Prog Neurobiol* 14:69-97

Moll R, Barto AG, Perkins TJ, Sutton RS. 1999. Learning instance-independent value functions to enhance local search. *Advances in Neural Information Processing Systems 11* 11:1017-23

90

Monchi O, Petrides M, Doyon J, Postuma RB, Worsley K, Dagher A. 2004. Neural bases of set-shifting deficits in Parkinson's disease. *J Neurosci* 24:702-10

Montague PR, Hyman SE, Cohen JD. 2004. Computational roles for dopamine in behavioural control. *Nature* 431:760-7

Montgomery EB, Jr., Nuessen J. 1990. The movement speed/accuracy operator in Parkinson's disease. *Neurology* 40:269-72

Morrish PK, Sawle GV, Brooks DJ. 1995. Clinical and [18F] dopa PET findings in early Parkinson's disease. *J Neurol Neurosurg Psychiatry* 59:597-600

Moustafa AA, Cohen MX, Sherman SJ, Frank MJ. 2008. A role for dopamine in temporal decision making and reward maximization in parkinsonism. *J Neurosci* 28:12294-304

Myers CE, Shohamy D, Gluck MA, Grossman S, Kluger A, et al. 2003. Dissociating hippocampal versus basal ganglia contributions to learning and transfer. *J Cogn Neurosci* 15:185-93

Naidich TP, Duvernoy, H.M., Delman, B.N., Sorensen, A.G., Kollias, S.S., Haacke, E.M. 2009. *Duvernoy's Atlas of the Human Brain Stem and Cerebellum*

NewYork: SpringerWien

Nair-Roberts RG, Chatelain-Badie SD, Benson E, White-Cooper H, Bolam JP, Ungless MA. 2008. Stereological estimates of dopaminergic, GABAergic and glutamatergic neurons in the ventral tegmental area, substantia nigra and retrorubral field in the rat. *Neuroscience* 152:1024-31

Nicholson SL, Brotchie JM. 2002. 5-hydroxytryptamine (5-HT, serotonin) and Parkinson's disease - opportunities for novel therapeutics to reduce the problems of levodopa therapy. *Eur J Neurol* 9 Suppl 3:1-6

Nicola SM. 2007. The nucleus accumbens as part of a basal ganglia action selection circuit. *Psychopharmacology (Berl)* 191:521-50

Niv Y. 2007. Cost, benefit, tonic, phasic: what do response rates tell us about dopamine and motivation? *Ann N Y Acad Sci* 1104:357-76

Niv Y, Daw ND, Joel D, Dayan P. 2007. Tonic dopamine: opportunity costs and the control of response vigor. *Psychopharmacology (Berl)* 191:507-20

Niv Y, Rivlin-Etzion M. 2007. Parkinson's disease: fighting the will? *J Neurosci* 27:11777-9

O'Doherty J, Dayan P, Schultz J, Deichmann R, Friston K, Dolan RJ. 2004. Dissociable roles of ventral and dorsal striatum in instrumental conditioning. *Science* 304:452-4

O'Doherty JP. 2004. Reward representations and reward-related learning in the human brain: insights from neuroimaging. *Curr Opin Neurobiol* 14:769-76

O'Doherty JP, Dayan P, Friston K, Critchley H, Dolan RJ. 2003. Temporal difference models and reward-related learning in the human brain. *Neuron* 38:329-37

O'Doherty JP, Hampton A, Kim H. 2007. Model-based fMRI and its application to reward learning and decision making. *Ann N Y Acad Sci* 1104:35-53

Obeso JA, Marin C, Rodriguez-Oroz C, Blesa J, Benitez-Temino B, et al. 2008. The basal ganglia in Parkinson's disease: current concepts and unexplained observations. *Ann Neurol* 64 Suppl 2:S30-46

Ogawa S, Lee TM, Kay AR, Tank DW. 1990a. Brain magnetic resonance imaging with contrast dependent on blood oxygenation. *Proc Natl Acad Sci U S A* 87:9868-72

Ogawa S, Lee TM, Nayak AS, Glynn P. 1990b. Oxygenation-sensitive contrast in magnetic resonance image of rodent brain at high magnetic fields. *Magn Reson Med* 14:68-78

Oullier O, Kelso JA. 2006. Neuroeconomics and the metastable brain. *Trends Cogn Sci* 10:353-4

Owen AM, James M, Leigh PN, Summers BA, Marsden CD, et al. 1992. Fronto-striatal cognitive deficits at different stages of Parkinson's disease. *Brain* 115 (Pt 6):1727-51

Padoa-Schioppa C, Assad JA. 2006. Neurons in the orbitofrontal cortex encode economic value. *Nature* 441:223-6

Pahwa R, Lyons K, Koller W, eds. 2003. *Handbook of Parkinson's disease*: Marcel Dekker

Palmiter RD. 2008. Dopamine signaling in the dorsal striatum is essential for motivated behaviors: lessons from dopamine-deficient mice. *Ann N Y Acad Sci* 1129:35-46

Parkinson JA, Dalley JW, Cardinal RN, Bamford A, Fehnert B, et al. 2002. Nucleus accumbens dopamine depletion impairs both acquisition and performance of appetitive Pavlovian approach behaviour: implications for mesoaccumbens dopamine function. *Behav Brain Res* 137:149-63

Pasupathy A, Miller EK. 2005. Different time courses of learning-related activity in the prefrontal cortex and striatum. *Nature* 433:873-6

Patterson CM, Newman JP. 1993. Reflectivity and learning from aversive events: toward a psychological mechanism for the syndromes of disinhibition. *Psychol Rev* 100:716-36

Pavese N, Brooks DJ. 2009. Imaging neurodegeneration in Parkinson's disease. *Biochim Biophys Acta* 1792:722-9

Pessiglione M, Seymour B, Flandin G, Dolan RJ, Frith CD. 2006. Dopamine-dependent prediction errors underpin reward-seeking behaviour in humans. *Nature* 442:1042-5

Plassmann H, O'Doherty JP, Rangel A. 2010. Appetitive and aversive goal values are encoded in the medial orbitofrontal cortex at the time of decision making. *J Neurosci* 30:10799-808

Praamstra P, Stegeman DF, Cools AR, Horstink MW. 1998. Reliance on external cues for movement initiation in Parkinson's disease. Evidence from movement-related potentials. *Brain* 121 (Pt 1):167-77

Pycock CJ, Kerwin RW, Carter CJ. 1980. Effect of lesion of cortical dopamine terminals on subcortical dopamine receptors in rats. *Nature* 286:74-6

Rabinovich MI, Huerta R, Varona P, Afraimovich VS. 2008. Transient cognitive dynamics, metastability, and decision making. *PLoS Comput Biol* 4:e1000072

Rahman S, Griffin HJ, Quinn NP, Jahanshahi M. 2008. The factors that induce or overcome freezing of gait in Parkinson's disease. *Behav Neurol* 19:127-36

Raz N, Rodrigue KM. 2006. Differential aging of the brain: patterns, cognitive correlates and modifiers. *Neurosci Biobehav Rev* 30:730-48

Redgrave P, Rodriguez M, Smith Y, Rodriguez-Oroz MC, Lehericy S, et al. 2010. Goal-directed and habitual control in the basal ganglia: implications for Parkinson's disease. *Nat Rev Neurosci* 11:760-72

Reynolds JN, Hyland BI, Wickens JR. 2001. A cellular mechanism of reward-related learning. *Nature* 413:67-70

Robbins TW. 2000. Chemical neuromodulation of frontal-executive functions in humans and other animals. *Exp Brain Res* 133:130-8

Robinson S, Sandstrom SM, Denenberg VH, Palmiter RD. 2005. Distinguishing whether dopamine regulates liking, wanting, and/or learning about rewards. *Behav Neurosci* 119:5-15

Roitman MF, Wheeler RA, Carelli RM. 2005. Nucleus accumbens neurons are innately tuned for rewarding and aversive taste stimuli, encode their predictors, and are linked to motor output. *Neuron* 45:587-97

Rowe JB, Hughes L, Ghosh BC, Eckstein D, Williams-Gray CH, et al. 2008. Parkinson's disease and dopaminergic therapy--differential effects on movement, reward and cognition. *Brain* 131:2094-105

Rutledge RB, Lazzaro SC, Lau B, Myers CE, Gluck MA, Glimcher PW. 2009. Dopaminergic drugs modulate learning rates and perseveration in Parkinson's patients in a dynamic foraging task. *J Neurosci* 29:15104-14

Salamone JD, Correa M. 2002. Motivational views of reinforcement: implications for understanding the behavioral functions of nucleus accumbens dopamine. *Behav Brain Res* 137:3-25

Salamone JD, Correa M, Mingote S, Weber SM. 2003. Nucleus accumbens dopamine and the regulation of effort in food-seeking behavior: implications for studies of natural motivation, psychiatry, and drug abuse. *J Pharmacol Exp Ther* 305:1-8

Salamone JD, Correa M, Mingote SM, Weber SM. 2005. Beyond the reward hypothesis: alternative functions of nucleus accumbens dopamine. *Curr Opin Pharmacol* 5:34-41

Samejima K, Ueda Y, Doya K, Kimura M. 2005. Representation of action-specific reward values in the striatum. *Science* 310:1337-40

Sarty G. 2007. *Computing Brain activity maps from fMRI time-series images*: Cambridge University Press

Sawaguchi T, Goldman-Rakic PS. 1991. D1 dopamine receptors in prefrontal cortex: involvement in working memory. *Science* 251:947-50

Schonberg T, O'Doherty JP, Joel D, Inzelberg R, Segev Y, Daw ND. 2010. Selective impairment of prediction error signaling in human dorsolateral but not ventral striatum in Parkinson's disease patients: evidence from a model-based fMRI study. *Neuroimage* 49:772-81

Schultz W. 1998. Predictive reward signal of dopamine neurons. *J Neurophysiol* 80:1-27

Schultz W. 2010. Dopamine signals for reward value and risk: basic and recent data. *Behav Brain Funct* 6:24

Schultz W, Dayan P, Montague PR. 1997. A neural substrate of prediction and reward. *Science* 275:1593-9

Sesack SR, Grace AA. 2010. Cortico-Basal Ganglia reward network: microcircuitry. *Neuropsychopharmacology* 35:27-47

Seymour B, McClure SM. 2008. Anchors, scales and the relative coding of value in the brain. *Curr Opin Neurobiol* 18:173-8

Sheridan MR, Flowers KA. 1990. Movement variability and bradykinesia in Parkinson's disease. *Brain* 113 (Pt 4):1149-61

Sheridan MR, Flowers KA, Hurrell J. 1987. Programming and execution of movement in Parkinson's disease. *Brain* 110 (Pt 5):1247-71

Shohamy D, Myers CE, Geghman KD, Sage J, Gluck MA. 2006. L-dopa impairs learning, but spares generalization, in Parkinson's disease. *Neuropsychologia* 44:774-84

Smith AD, Bolam JP. 1990. The neural network of the basal ganglia as revealed by the study of synaptic connections of identified neurones. *Trends Neurosci* 13:259-65

Smith SM. 2004. Overview of fMRI analysis. *Br J Radiol* 77 Spec No 2:S167-75

Smith Y, Bevan MD, Shink E, Bolam JP. 1998. Microcircuitry of the direct and indirect pathways of the basal ganglia. *Neuroscience* 86:353-87

Sokoloff P, Giros B, Martres MP, Bouthenet ML, Schwartz JC. 1990. Molecular cloning and characterization of a novel dopamine receptor (D3) as a target for neuroleptics. *Nature* 347:146-51

Sommer F, Wichert A, eds. 2003. *Exploratory Analysis and Data Modeling in Functional Neuroimaging*: Massachusetts Institute of Technology Press

Stuss DT, Knight RT, eds. 2002a. *Principals of frontal lobe function*: Oxford university press

Stuss DT, Knight RT, eds. 2002b. *Principles of frontal lobe function*: Oxford university press

Suhara T, Fukuda H, Inoue O, Itoh T, Suzuki K, et al. 1991. Age-related changes in human D1 dopamine receptors measured by positron emission tomography. *Psychopharmacology (Berl)* 103:41-5

Suteerawattananon M, Morris GS, Etnyre BR, Jankovic J, Protas EJ. 2004. Effects of visual and auditory cues on gait in individuals with Parkinson's disease. *J Neurol Sci* 219:63-9

Sutton RB, AG. 1998. *Reinforcement learning, An Introduction*. Cabridge, Massachusetts: MIT press

Swanson LW. 1982. The projections of the ventral tegmental area and adjacent regions: a combined fluorescent retrograde tracer and immunofluorescence study in the rat. *Brain Res Bull* 9:321-53

Teasdale N, Phillips J, Stelmach GE. 1990. Temporal movement control in patients with Parkinson's disease. *J Neurol Neurosurg Psychiatry* 53:862-8

Thulborn KR, Waterton JC, Matthews PM, Radda GK. 1982. Oxygenation dependence of the transverse relaxation time of water protons in whole blood at high field. *Biochim Biophys Acta* 714:265-70

Ungerstedt U. 1971. Striatal dopamine release after amphetamine or nerve degeneration revealed by rotational behaviour. *Acta Physiol Scand Suppl* 367:49-68

Van Bockstaele EJ, Cestari DM, Pickel VM. 1994. Synaptic structure and connectivity of serotonin terminals in the ventral tegmental area: potential sites for modulation of mesolimbic dopamine neurons. *Brain Res* 647:307-22

van Schouwenburg M, Aarts E, Cools R. 2010. Dopaminergic modulation of cognitive control: distinct roles for the prefrontal cortex and the basal ganglia. *Curr Pharm Des* 16:2026-32

Voon V, Pessiglione M, Brezing C, Gallea C, Fernandez HH, et al. 2010. Mechanisms underlying dopamine-mediated reward bias in compulsive behaviors. *Neuron* 65:135-42

Voon V, Sohr M, Lang AE, Potenza MN, Siderowf AD, et al. 2011. Impulse control disorders in Parkinson disease: a multicenter case--control study. *Ann Neurol* 69:986-96

Weintraub D, Hoops S, Shea JA, Lyons KE, Pahwa R, et al. 2009. Validation of the questionnaire for impulsive-compulsive disorders in Parkinson's disease. *Mov Disord* 24:1461-7

Weintraub D, Koester J, Potenza MN, Siderowf AD, Stacy M, et al. 2010. Impulse control disorders in Parkinson disease: a cross-sectional study of 3090 patients. *Arch Neurol* 67:589-95

Wiesendanger E, Clarke S, Kraftsik R, Tardif E. 2004. Topography of cortico-striatal connections in man: anatomical evidence for parallel organization. *Eur J Neurosci* 20:1915-22

Williams-Gray CH, Evans JR, Goris A, Foltynie T, Ban M, et al. 2009. The distinct cognitive syndromes of Parkinson's disease: 5 year follow-up of the CamPaIGN cohort. *Brain* 132:2958-69

Williams GV, Goldman-Rakic PS. 1995. Modulation of memory fields by dopamine D1 receptors in prefrontal cortex. *Nature* 376:572-5

Wilson JM, Levey AI, Rajput A, Ang L, Guttman M, et al. 1996. Differential changes in neurochemical markers of striatal dopamine nerve terminals in idiopathic Parkinson's disease. *Neurology* 47:718-26

Wise RA. 1978. Catecholamine theories of reward: a critical review. *Brain Res* 152:215-47

Wise RA. 2004. Dopamine, learning and motivation. *Nat Rev Neurosci* 5:483-94

Wise RA, Rompre PP. 1989. Brain dopamine and reward. *Annu Rev Psychol* 40:191-225

Wolf ME, Sun X, Mangiavacchi S, Chao SZ. 2004. Psychomotor stimulants and neuronal plasticity. *Neuropharmacology* 47 Suppl 1:61-79

Wu M, Brudzynski SM. 1995. Mesolimbic dopamine terminals and locomotor activity induced from the subiculum. *Neuroreport* 6:1601-4

Wunderlich K, Rangel A, O'Doherty JP. 2009. Neural computations underlying action-based decision making in the human brain. *Proc Natl Acad Sci U S A* 106:17199-204

Wunderlich K, Rangel A, O'Doherty JP. 2010. Economic choices can be made using only stimulus values. *Proc Natl Acad Sci U S A* 107:15005-10

Yin HH, Ostlund SB, Balleine BW. 2008. Reward-guided learning beyond dopamine in the nucleus accumbens: the integrative functions of cortico-basal ganglia networks. *Eur J Neurosci* 28:1437-48

Yin HH, Zhuang X, Balleine BW. 2006. Instrumental learning in hyperdopaminergic mice. *Neurobiol Learn Mem* 85:283-8