Dynamic Dopamine Modulation of Striato-Cortical Circuits in Cognition: Converging Neuropsychological, Psychopharmacological and Computational Studies

by

Michael Joshua Frank
B.S., Queen's University, 1997
M.S., University of Colorado, 2000

A thesis submitted to the
Faculty of the Graduate School of the
University of Colorado in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
Department of Psychology and Center for Neuroscience
2004

This thesis entitled:
Dynamic Dopamine Modulation of Striato-Cortical Circuits in Cognition: Converging Neuropsychological, Psychopharmacological and Computational Studies
written by Michael Joshua Frank
has been approved for the Department of Psychology and Center for Neuroscience

Randall C. O'Reilly

Tim Curran

Jerry W. Rudy

Date

The final copy of this thesis has been examined by the signatories, and we find that both the content and the form meet acceptable presentation standards of scholarly work in the above mentioned discipline.

Frank, Michael Joshua (Ph.D., Psychology and Neuroscience)

Dynamic Dopamine Modulation of Striato-Cortical Circuits in Cognition: Converging Neuropsychological, Psychopharmacological and Computational Studies

Thesis directed by Dr. Randall C. O'Reilly

How do we produce complex motor sequences? To what extent do we learn from the positive versus negative consequences of our decisions? How do we maintain task-relevant information in working memory while ignoring distracting information? This dissertation provides a mechanistic framework that explores how these seemingly unrelated processes recruit remarkably similar neural circuits linking the basal ganglia (BG) with frontal cortex. Drawing from neuroanatomical and biochemical considerations, this framework suggests that the BG facilitate or suppress cortical "actions" (e.g., motor responses and working memory updating) via separate Go and NoGo pathways projecting to frontal cortex, and that the relative balance of these pathways is dynamically modulated by dopamine (DA). Transient DA bursts and dips during positive and negative reinforcement support Go and NoGo learning via D1 and D2 receptors, respectively. Computational neural network models instantiate key biological properties and provide insight into the underlying role of BG/DA interactions during the learning and execution of cognitive tasks. These models account for complex medication-dependent cognitive deficits in Parkinson's disease, and make simple predictions for the underlying source of these deficits, emphasizing the importance of the dynamic range of DA signals. These predictions have been subsequently confirmed in medicated and non-medicated Parkinson's patients and in healthy individuals under pharmacologically-induced DA manipulation. In all of these studies, elevated levels of phasic DA release led to greater Go learning from positive outcomes of decisions, whereas diminished DA levels led to better NoGo learning to avoid negative outcomes. Tonic DA stimulation led to more overall Go responding. These effects extended to higher level cognitive function: tonic DA stimulation led to more overall working memory updating and concomitant distractibility, whereas enhanced phasic DA release led to greater selective updating for task-relevant (i.e., "positively-valenced") information, but difficulty in ignoring this information in a subsequent set-shift. Drug effects also interacted with baseline working memory span. Taken together, these results provide substantial support for a unified account of the role of DA in modulating cognitive processes that depend on the basal ganglia.

Acknowledgements

This dissertation would have not been possible without the support of many people and institutions throughout the last several years. First and foremost, I want to thank my supervisor Dr. Randall O'Reilly for being a tremendous mentor and friend. Randy has been by far the most influential person in all of my years of education. I am also grateful to the members of my dissertation committee for their guidance and support. Over the years, Drs. Marie Banich, Tim Curran and Jerry Rudy have all provided me with exceptional training in different but complementary ways. I am indebted to Dr. Yuko Munakata, who has engaged me in several interesting discussions and has provided me with invaluable guidance, support, and friendship. I also want to thank my fellow graduate students in the cognitive psychology and neuroscience programs, as well as other members of the CCN lab, for thought-provoking discussion about research and bigger picture issues.

For my research with Parkinson's patients, I gratefully acknowledge Dr. Seeberger and Vicki Segro and staff members of the Colorado Neurological Institute, who were able to support our project despite their busy and important jobs in medical care. For the psychopharmacological studies, I thank Dr. Perreault, Emily Edwards, David Bynum, Lulu Davis and Sue Martin and other helpful and supportive staff members at the General Clinical Research Center.

Finally, I am very thankful to all my family and friends for their tremendous support throughout all my endeavors. Sadly, my grandfather Gil Frank passed away a week before my defense. He was a self-educated man who was active in many intellectual, inter-cultural and environmental issues. To his memory I dedicate these pages.

Contents

Chapter

List of Tables

Table

List of Figures

Figure

Chapter 1

Introduction

The capacity for executive function (EF) gives us the power to dictate our own actions based on their perceived consequences, and lies at the core of what makes us human. Other animals also show signs of EF, albeit to a lesser extent. A primate's ability to consider alternatives before making decisions, to override prepotent responses to achieve longer-term satisfaction, and to flexibly switch focuses of attention likely form the foundations for EF in humans. The neural substrates that underly these primitives may also support complex problem solving and the tendency to ruminate about decisions to guide future plans. Once considered beyond the grasp of cognitive neuroscience, the neurobiological basis for EF is now within our reach.

Several years of research has demonstrated that across both phylogeny and ontogeny, EF emerges as a function of evolution and development of the prefrontal cortex (PFC) and dopaminergic systems. The PFC allows animals to "hold ideas in mind" by actively maintaining recent information via persistent neural firing (Miller et al., 1996; Fuster, 1997; Goldman-Rakic, 1995). It supports goal-directed behavior by providing top-down support to guide motor systems (e.g., Cohen et al., 1990). Dopamine (DA) in the PFC keeps the system "well-greased" by enhancing the maintenance of task-relevant information while making this activity less susceptible to distracting interference (Williams & Goldman-Rakic, 1995; Goldman-Rakic, 1996; Durstewitz et al., 2000a; Dreher & Burnod, 2002). But while the maintenance of information over time may be a satisfactory account of short term memory, it falls short of being able to explain complexities associated with more advanced forms of EF, such as flexible shifting between attentional sets. In particular, three key questions are (a) how does the system differentiate between relevant information to be stored in active form and distracting information that should be ignored? (b) how do modulatory properties of DA that change dynamically in response to positive and negative reinforcement play a role in these processes? and (c) what is the function of subcortical areas that interact with PFC?

Ironically, some insight into these complex issues may come from studying the more primitive motor system. To produce smooth movements, a "controller" must dictate which motor commands to execute and which to suppress at a given instant. Analogously, for cohererent thought, a prefrontal controller must dictate which information to actively maintain over time and which information to ignore. But the analogy is not one of mere functional similarity. It is generally accepted that the basal ganglia (BG) act as the motor controller by dynamically modulating activity in frontal motor cortex (e.g., Hikosaka, 1998). Similarly, various researchers now propose a key role of parallel circuits linking the BG, thalamus, and PFC that are astonishingly similar to those involved in the motor circuit (Alexander, DeLong, & Strick, 1986; Casey, Durston, & Fosella, 2001a; Middleton & Strick, 2002). But a unified account of **how** the BG contributes to frontal function, and what the critical roles of DA are in these regions, has yet to be explored. Given that dysfunction in these circuits (and not just PFC itself) is thought to give rise to cognitive deficits in neurological disorders ranging from Parkinson's disease, obsessive compulsive disorder, attention deficit/hyperactivity disorder, schizophrenia, and Tourrette's syndrome, among others (Nieoullon, 2002; Mink, 2001), a more explicit theoretical framework for establishing the role of the BG in higher level cognition is in order.

In this dissertation I argue that the fundamental role of the BG is to provide "Go" and "NoGo" signals which facilitate or suppress the execution of "actions" being considered in frontal cortex, and that the type of

action depends on the cortical region of interest. Dopamine modulates the relative balance of these BG signals, effectively modulating the threshold at which actions are facilitated versus suppressed. This effect is dynamic, such that transient changes in DA levels during positive and negative reinforcement support the learning of Go and NoGo signals that facilitate adaptive actions while suppressing nonadaptive actions. I argue that this functionality is very similar across the domains of motor control (learning adaptive versus nonadaptive motor sequences), decision making (making choices that are more likely to result in positive rather than negative outcomes), and executive function (paying attention to task-relevant versus irrelevant information).

Note that with respect to executive function, I am not suggesting that the BG/DA system in fact accomplishes what is traditionally thought to be subserved by PFC. Rather, I argue that DA processes within the BG play a complementary role to maintenance-enhancing functions of DA in PFC by modulating the degree to which information in PFC gets updated. The theory is leveraged from existing knowledge of BG involvement in motor control, in which it is thought to selectively facilitate a preferred motor command in premotor cortex while suppressing competing motor programs (Mink, 1996; Jiang et al., 2003). In PFC, "commands" to be facilitated are not motor but are instead whether or not to store current sensory information in working memory. Thus the BG acts as a gate on PFC working memory representations: a BG Go signal results in updating and subsequent maintenance of PFC activity, whereas a NoGo signal prevents the PFC from maintaining current sensory information (while it continues to maintain previous working memory representations). Transient increases in DA result in more Go and less NoGo activity, providing a **dynamic gating signal** on frontal representations (Frank et al., 2001; Frank, in press). In the context of EF, this gating signal is necessary to allow the PFC to flexibly and rapidly shift between tasks and attentional sets (Braver & Cohen, 2000; O'Reilly et al., 2002; Rougier & O'Reilly, 2002; Rougier et al., submitted).

Countless articles have been published regarding the important neurotransmitter dopamine: a recent citation search with the single term "dopamine" listed 91,894 results. Again though, systems-level accounts of DA function in cognition are rare. For instance, many emphasize the critical role of DA in the PFC in executive function. But while D1 receptors are abundant in PFC and are critical for EF, there is ample evidence that effects of D2 receptor stimulation are more pronounced in the BG (Camps et al., 1989; Goldman-Rakic et al., 1990; Arnsten et al., 1995). These results can only be interpreted in the context of a theoretical understanding of the dynamics of striato-cortical circuits. This dissertation takes advantage of the wealth of DA findings in reward, learning, motivation, and motor control by contextualizing them in the background of BG-PFC circuits. What emerges is a theoretical framework for understanding the role(s) of DA in the BG and PFC in higher level cognition, which can then be tested experimentally. To get there, I will summarize various aspects of the literature, integrating converging evidence across multiple levels of analysis, and drawing from computational modeling of systems-level interactions between BG and PFC.

1.0.1 Overview and Framework

By focusing on functional effects of DA in frontal cortex (FC) and basal ganglia (BG) in the context of tasks that depend on these regions, the first part of the dissertation establishes a framework for understanding how these cortical and subcortical areas interact to support cognitive reinforcement learning and executive function processes. The general layout of this framework is summarized below.

- The BG acts as a **modulator** by suppressing behaviors that are being "considered" in cortex and only enabling their execution when conditions are appropriate. Two main populations of cells in the striatum compete to facilitate or suppress these behaviors. Mediated by parallel BG-FC circuits, the type of behavior modulated is distinct in different parts of frontal cortex — the BG suppresses/facilitates motor responses in motor cortex, and may have analogous functions for updating working memory/attentional information in PFC, and motivational/reward information in orbitofrontal cortex (OFC).

- Dopamine in the BG modulates activity in an already modulatory system, by enhancing neural activity of BG cells that facilitate behaviors, while concurrently inhibiting those that would suppress behaviors.

This effect is dynamic, such that increases in extracellular DA during positive reinforcement drive Go learning, whereas decreases in DA during negative reinforcement drive NoGo learning.

- DA neurons fire both **tonically** and **phasically**. Tonic DA affects performance (elevated levels of DA are associated with increased responsiveness and more PFC updating, whereas decreased levels of DA result in hesitancy and less PFC updating). Phasic DA firing is related to reward and expectations thereof, and directly drive learning. Furthermore, various pharmacological and clinical effects on the DA system can differentially affect tonic and phasic firing, which can explain observed cognitive effects.

- In PFC, a healthy amount of DA acts to stabilize memory representations, supporting robust maintenance of information over time and in the face of distracting interference.

This framework demonstrates that "frontal" impairments in various neurological conditions may not only arise from damage to PFC itself, but can also result from dopaminergic dysfunction within the BG. A considerable amount of cognitive behavioral, neuropsychological and pharmacological observations can be accounted for by this model, including some that are counterintuitive when PFC is only considered by itself.

This dissertation is organized as follows. The remainder of this chapter begins with a review of cellular and systems-level effects of DA in the BG and PFC (section 1.1). Because these effects strongly depend on the level of DA, section 1.2 discusses behavioral events (mainly concerning reward) that cause dynamic bursting and dips in DA release, implicating a role for DA in motivation and addiction. Section 1.3 suggests that the fairly well known role of BG DA in motor systems can be extended to cognitive processes, forming the basis for the primary hypotheses of the paper. Correlations between development/decline of the DA system and cognitive function are then introduced, but are in and of themselves inconclusive. Section 1.5 provides converging evidence for the framework by examining cognitive effects resulting from selective DA lesions in animals, while section 1.6 discusses possible higher-level implications for patients with ADHD. Following this theoretical overview are the use of computational models to explore the framework's potential for understanding cognitive reinforcement learning deficits in Parkinson's patients (chapter 2). The model is then extended to include higher level interactions in working memory and executive function (chapter 3). Next, behavioral experiments provide confirmation of simple model predictions for the underlying source of medication-dependent cognitive deficits in Parkinson's disease (chapter 4). Finally, chapter 5 describes a rigorous test of the framework across a wide range of cognitive tasks in healthy individuals under pharmacological challenge.

1.1 The Basal Ganglia as a Gate on Cortical Function

In the context of motor control, various authors have suggested that the BG selectively facilitates the execution of a single motor command, while suppressing all others (Mink, 1996; Chevalier & Deniau, 1990; Frank et al., 2001; Frank, in press; Jiang et al., 2003). Thus, the BG is thought to act as a brake on competing motor actions that are represented in motor cortex — only the most appropriate motor command is able to "release the brake" and get executed at any point in time. Further, the BG does not come up with the motor responses itself, but instead modulates the execution of cortical responses by signaling "Go" or "NoGo" (Hikosaka, 1989). Phasic changes in dopamine during reinforcement (or lack thereof) modulate Go/NoGo signals which support implicit learning of stimulus-response mappings (Frank, in press). This functionality also helps to string simple motor commands together to form a complex motor sequence, by selecting the most appropriate command at any given portion of the sequence and inhibiting the other ones until the time is appropriate (Mink, 1996). The circuitry that implements these functions is described next.

The BG consists of a number of structures which work together as a system (figures 1.1 and 1.2). The input and largest segment of the BG is generically termed "striatum", which is further divided into the caudate nucleus, putamen, and ventral striatum/nucleus accumbens. Each striatal area receives and modulates function in distinct parts of cortex that are involved in multiple facets of cognition. The putamen receives from and

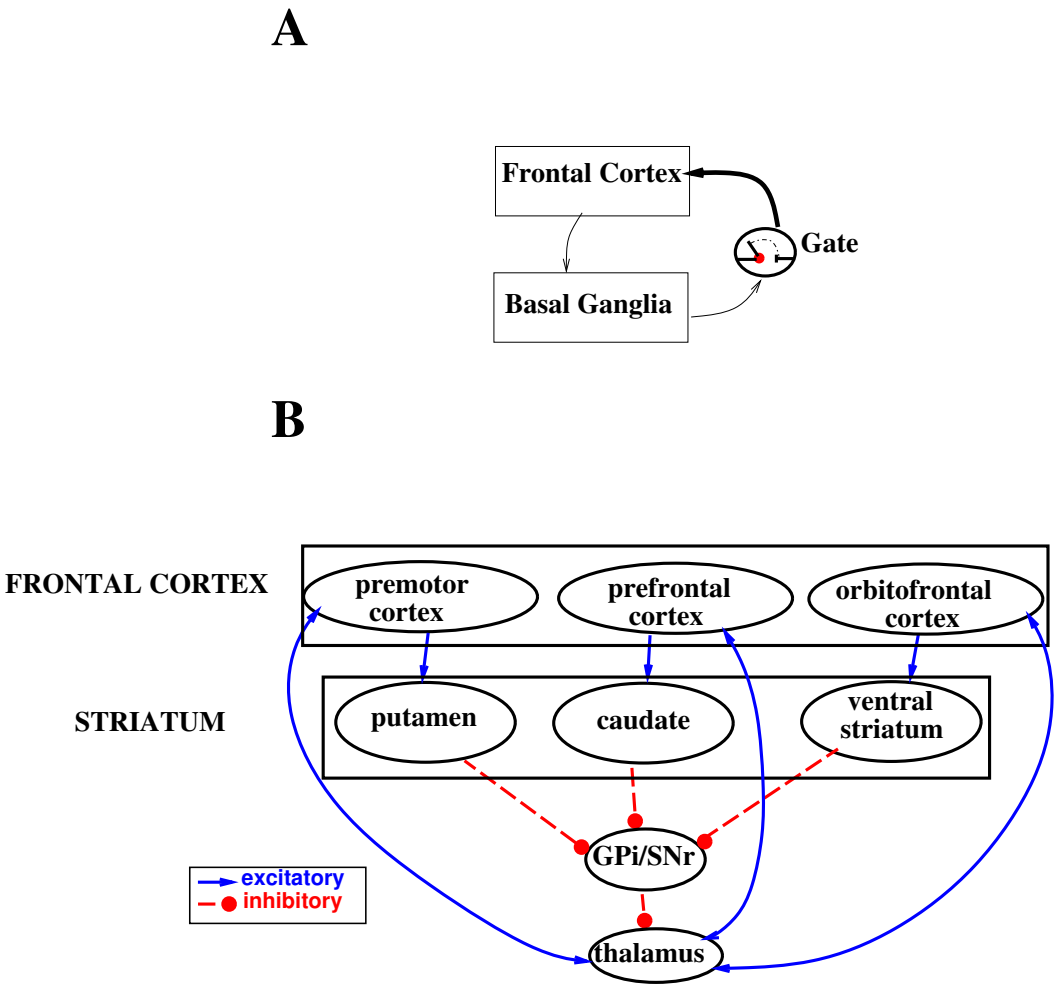

Figure 1.1: A) The basal ganglia acts as a gate on frontal cortical representations, by receiving information about its current plans and either suppressing or facilitating them. B) BG-FC circuits work together in many facets of cognition. These include motor planning in premotor cortex, working memory and cognitive control in dorsolateral PFC, and reward information in orbitofrontal cortex. The throughput for the gating function is the thalamus, which, when excited, enhances cortical representations via recurrent thalamocortical excitatory connectivity. The BG normally suppresses, or acts as a brake on cortical representations, because its output structures (globus pallidus and substantia nigra) tonically inhibit the thalamus. To facilitate a cortical command, striatal neurons become excited which send inhibitory projections to GP/SNr, thereby **disinhibiting** the thalamus. This disinhibition does not itself excite the thalamus, but only enables it to get excited from descending cortical projections. Thus the BG gates cortical representations by "releasing its brake" on them. Shown here are three of the five BG-FC loops initially described in Alexander et al (1986).

modulates premotor cortex, the caudate nucleus – prefrontal cortex, and the ventral striatum – orbitofrontal cortex. Aside from differences in anatomical location and cortical interconnectivity, neurons in these striatal areas have very similar characteristics. The vast majority (90-95%) of all striatal neurons are GABAergic medium spiny neurons (Gerfen & Wilson, 1996). These are projection cells that carry information through BG output structures (globus pallidus and substantia nigra) to the thalamus, ultimately closing the circuit back to the area of cortex from which they received (Alexander et al., 1986; Alexander & Crutcher, 1990a).

To facilitate/suppress a cortical command, the BG can either excite or inhibit areas of the thalamus that provide bottom-up excitation to cortex. Two main pathways in the BG are thought to provide these two functions (Alexander & Crutcher, 1990a). The "direct" pathway facilitates the execution of responses (Go), whereas the "indirect" pathway inhibits them (NoGo). Cells in the direct pathway project from the striatum and inhibit the internal segment of the globus pallidus (GPi).[1] In the absence of striatal firing, the GPi tonically inhibits the thalamus, so direct pathway activity and resulting GPi inhibition serves to **disinhibit** the thalamus. The double-negative invoked by this disinhibition provides a gating function on cortical commands: it does not directly excite the thalamus but instead simply **enables** the thalamus to get excited from other excitatory projections (e.g., Chevalier & Deniau, 1990; Frank, Loughry, & O'Reilly, 2001, figure 1.1). Cells in the indirect pathway inhibit the external segment of the globus pallidus (GPe), which tonically inhibits the GPi.[2] The net effect of indirect pathway activity is then to further inhibit the thalamus and suppress the cortical command from getting executed. See figure 1.2 for a pictorial description of this circuitry.

1.1.1 Cellular and Systems-Level Mechanisms of DA in the BG

1.1.1.1 But First, Some Caveats

Because an intact dopaminergic system is integral for a wide variety of cognitive functions, it would be useful to outline the mechanism by which DA acts on neurons in brain regions that are substrates for these functions. Unfortunately, because dopamine is a neuromodulator (i.e., its release is tonic and diffuse), its post-synaptic effects are not straightforwardly characterized. Despite the enormous amount of research conducted on this topic, the data are overwhelmingly complex and contradictory, and consequently there is disagreement regarding not only systems-level function of dopamine release, but even its immediate cellular effects.

The mechanisms by which dopamine increases or decreases the excitability of neurons has been shown to involve modulation of sodium, potassium and calcium currents (Nicola et al., 2000). Unfortunately for those attempting to consider broader implications, these effects can vary (a) in different parts of the brain; (b) in studies performed **in vitro** versus those performed **in vivo**; (c) in different behavioral conditions, and (d) depending on interactions with other neurotransmitters such as acetylcholine, norepinephrine, serotonin, glutamate, and GABA.

Below I will focus on dopaminergic effects on just D1 and D2 receptors, not only because these are the most well studied and understood but also because these are the two general classes of receptors. The D1 class of DA receptors includes D1 and D5, whereas the D2 class includes D2, D3, and D4. The assignment of these receptors to the two general classes is based on the type of G protein activated by the receptor and whether it stimulates (D1 class) or inhibits (D2 class) adenylyl cyclase (Nicola et al., 2000). These molecular effects ultimately lead to opposite roles of DA in modulating synaptic plasticity in the two classes of receptors.

It should also be kept in mind that DA is but a single neurotransmitter which by itself cannot implement any behavioral process, cognitive or otherwise. As such, its effects cannot be fully appreciated without first contextualizing the functional circuitry of the brain region it is modulating. I now turn to a general (if highly simplified) description of BG circuitry and function. I will then review the cellular effects of DA in the BG and briefly describe those in frontal cortex.

[1] The substantia nigra pars reticulata (SNr) is equivalent to the GPi in this circuitry, except that the former receives from the caudate, and the latter from the putamen. For this reason I consider the two to be one functional entity, but for simplicity only refer to GPi.

[2] The GPe inhibition of GPi also involves the subthalamic nucleus, but is left out of this discussion for simplicity.

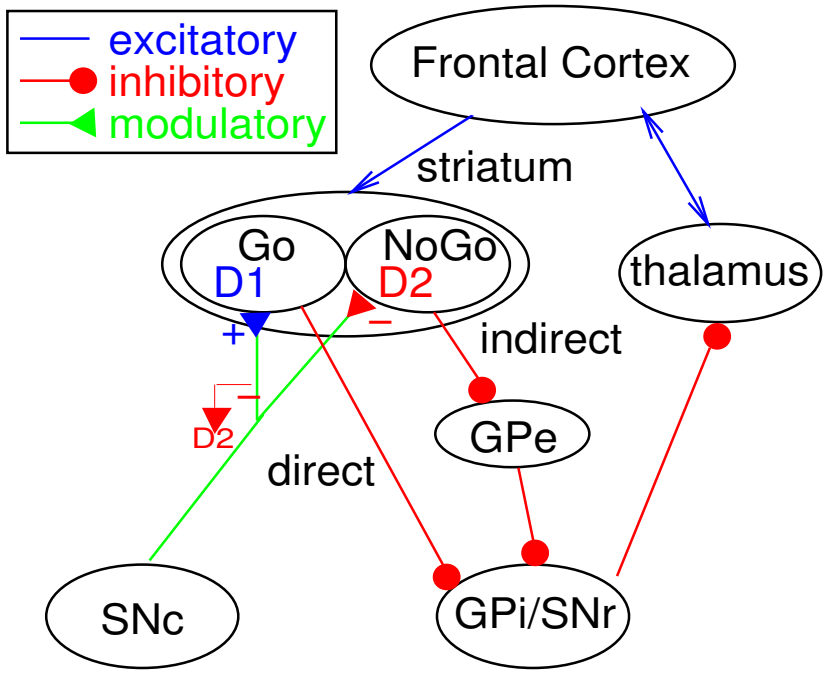

Figure 1.2: The cortico-striato-thalamo-cortical loops, including the direct and indirect pathways of the basal ganglia. The cells of the striatum are divided into two sub-classes based on differences in biochemistry and efferent projections. The "Go" cells project directly to the GPi, and have the effect of disinhibiting the thalamus, thereby facilitating the execution of an action represented in cortex. The "NoGo" cells are part of the indirect pathway to the GPi, and have an opposing effect, suppressing actions from getting executed. Dopamine from the SNc projects to the dorsal striatum, differentially modulating activity in the direct and indirect pathways by activating different receptors: The Go cells express the D1 receptor, and the NoGo cells express the D2 receptor. D2 receptors are also present presynaptically, and provide negative feedback on the amount of DA release during bursting. Dopamine from the VTA projects to ventral striatum and frontal cortex (not shown). GPi: internal segment of globus pallidus; GPe: external segment of globus pallidus; SNc: substantia nigra pars compacta; SNr: substantia nigra pars reticulata.

1.1.1.2 Cellular and Network Effects of BG/DA

The systems-level effect of DA in the BG is to dynamically modulate the relative balance of Go/NoGo signals. This is because (a) cells in the direct pathway predominantly express the D1 class of DA receptors, whereas those in the indirect pathway predominantly express the D2 class (Gerfen, 1992; Gerfen & Keefe, 1994; Bloch & LeMoine, 1994; Le Moine & Bloch, 1995; Gerfen et al., 1995; Ince et al., 1997; Aubert et al., 2000); and (b) dopamine is primarily excitatory to cells expressing D1 receptors but is inhibitory to those expressing D2 receptors (e.g., Hernandez-Lopez et al., 1997; Hernandez-Lopez et al., 2000). Increased levels of DA enhance activity in the direct/Go pathway while concurrently suppressing the indirect/NoGo pathway (e.g., Gerfen & Wilson, 1996; Gurney et al., 2001; Akkal et al., 1996). Decreased levels of DA have the opposite effect, releasing the indirect pathway from suppression. As such, chronic DA depletion of the striatum (as in Parkinson's disease) causes the indirect pathway to be tonically overactive (Gerfen, 2000; Salin et al., 1996). Finally, DA levels are under tight feedback control by autoreceptors (of the D2 type) acting to limit the amount of phasic DA released during bursting. These effects (summarized in Table 1.1) form the basis for interpreting the role of DA in cognitive processes and neurological disorders described below. First, I discuss the relevant data in more detail.

1.1.1.3 D1 and D2 Receptors are Segregated in Direct and Indirect Pathways

D1 and D2 receptors are relatively segregated in the direct and indirect pathways, a consistent observation that has been confirmed in both rats and primates (Gerfen, 1992; Aubert, Ghorayeb, Normand, & Bloch, 2000). While some caution that D1/D2 receptors are co-localized in cells within both pathways (Aizman et al., 2000; Surmeier et al., 1996), they reconcile these conflicting data by noting that the relative levels of expression are asymmetrical. That this asymmetry is functionally significant is evident by observations that D1 and D2 agents induce immediate early gene expression (a marker of neuronal activity) in spiny projection neurons in the direct and indirect pathways, respectively (Robertson et al., 1992; Gerfen et al., 1995). These observations are suggestive of a mechanism for differential DA modulation of Go and NoGo signals, which has also been observed with neuroimaging techniques in humans under pharmacological DA manipulation (Black et al., 1997). Other data shows that co-activation of both striatal DA receptors is sometimes necessary for controlling BG output (e.g., Waszczak et al., 2002). This is consistent with the idea that successful BG function requires the simultaneous activation of an appropriate Go response (supported by D1 activation) with inhibition of its NoGo representation (supported by D2 activation).

1.1.1.4 D1 Enhances Go Activity

That dopamine is a modulator, and not a simple neurotransmitter, is perhaps best exemplified by its effects on the D1 receptor. D1 stimulation can be either excitatory or inhibitory, depending on the state of the target cell. In vivo recordings demonstrate that spiny cells oscillate between two states of resting membrane potential (Wilson, 1993; Wilson & Kawaguchi, 1996; Cowan et al., 1994). Synaptic input from multiple afferents can induce the "up-state" (resting membrane potential around -50 mV), making the cell easily excitable (Blackwell et al., 2003; O'Donnell & Grace, 1995). Without sufficient synaptic drive, "down-state" cells (membrane potential around -80 mv) only fire spuriously. Interestingly, dopamine D1 stimulation further excites up-state cells but inhibits those in the down-state (Nicola et al., 2000; Hernandez-Lopez et al., 1997). At a molecular level, D1 activation enhances L-type Ca^{2+} channel currents in striatal medium spiny neurons (Surmeier et al., 1995), which mediate both the excitatory and inhibitory effects (Hernandez-Lopez et al., 1997).

These observations are consistent with the common notion that D1 stimulation enhances contrast or increases the signal-to-noise ratio (Rolls et al., 1984; Foote & Morrison, 1987; Cohen et al., 2002), because it amplifies incoming synaptic drive while suppressing spontaneous background "noise". In the BG, this may allow DA to enhance the Go signal for facilitating the most appropriate response while concurrently suppressing Go activity for competing responses (Frank, in press).

Condition	Go	NoGo	DA burst
D1 ↑ (stimulation, agonist)	+		
D1 ↓ (blockade, antagonist)	–		
D2 ↑ (stimulation, agonist)		–	–
D2 ↓ (blockade, antagonist)		+	+

Table 1.1: Summary of Go/NoGo effects of the D1 and D2 receptors, and presynaptic D2 effects on DA bursting.

This hypothesis also reconciles several seemingly incongruent results. Because cells studied **in vitro** are typically in their down-states (due to the absence of cortical input), previous **in vitro** work demonstrated an inhibitory effect of D1 activation. **In vivo**, cells oscillate between the up and down states, and in the context of this hypothesis it is not surprising that **in vivo** studies have reported both excitatory and inhibitory effects of D1 receptor activation (Kiyatkin & Rebec, 1996; Mercuri et al., 1985; Kitai et al., 1976).

1.1.1.5 D2 Inhibits NoGo Activity ("Releasing the Brakes")

Whereas D1 receptors are found abundantly in cortical regions, cells expressing the D2 receptors predominate in subcortical areas (Szele et al., 1991). The striatal D2 receptor density is 11 times greater than frontal cortex and 27 times greater than occipital cortex (Camps et al., 1989). As such, the effects of dopamine on striatal cells expressing the D2 receptor is likely to be critical. Further, systemic administration of D2 agonists/antagonists used to treat various neurological conditions are likely to act primarily and most directly in the striatum. D2 receptors are expressed both postsynaptically (like the D1 receptors) and as presynaptic autoreceptors that tightly control the level of phasic DA release via negative feedback (Starke et al., 1989; Schmitz et al., 2003). In both pre and postsynaptic cases, D2 activation is inhibitory (e.g., Hernandez-Lopez et al., 2000). Nevertheless, the functional consequences of D2 inhibition on pre and postsynaptic receptors are markedly different, as described next.

As noted above, postsynaptic D2 receptors are most prevalent in the indirect pathway, and D2 stimulation is inhibitory. Dopamine therefore has a net suppressive effect on indirect activity. Since this activity normally acts as a brake on cortical commands by signaling NoGo, increased levels of DA release the brake by inhibiting NoGo activity and allowing the Go pathway to exert more influence on BG output (Albin et al., 1989). By this account, dopamine shifts the balance in the BG from being "hesitant" to a more responsive state, effectively **lowering the threshold for facilitating the execution of cortical commands**. Parkinson's patients have difficulty initiating motor commands because a lack of BG DA results in an overactive indirect pathway leading to excessive cortical inhibition (Filion & Tremblay, 1991; Jellinger, 2002)—in other words, a tonic state of NoGo. Medication may restore DA function so that the balance is shifted to Go, allowing patients to execute responses more fluidly. Additional support for this account comes from the common observations that D2 agonists have a psychomotor stimulant effect on behavior, while D2 antagonists induce catalepsy and Parkinsonism (Fog, 1972), apparently by inhibiting and enhancing NoGo, respectively.

1.1.1.6 Presynaptic D2 Receptors Inhibit Phasic DA Release

Just as D2 stimulation is inhibitory on postsynaptic activity, it is also inhibitory on presynaptic D2 autoreceptors. However, because these autoreceptors control phasic DA release, stimulating them leads to inhibition of DA release, and therefore has opposing effects to postsynaptic D2 stimulation. For example, D2 agonists stimulate these autoreceptors and lead to decreased phasic DA release, whereas antagonists block autoreceptors and actually **increase** DA bursting and release (Wu et al., 2002; Garris et al., 2003) (Figure 5.2). Thus, taking both pre and postsynaptic effects into account, a D2 agonist will tonically suppress NoGo (lowering the threshold for action and therefore more Go responding) while also inhibiting phasic DA release (which impairs the learning of positive outcomes of these Go responses). In contrast, a D2 antagonist will tonically enhance NoGo

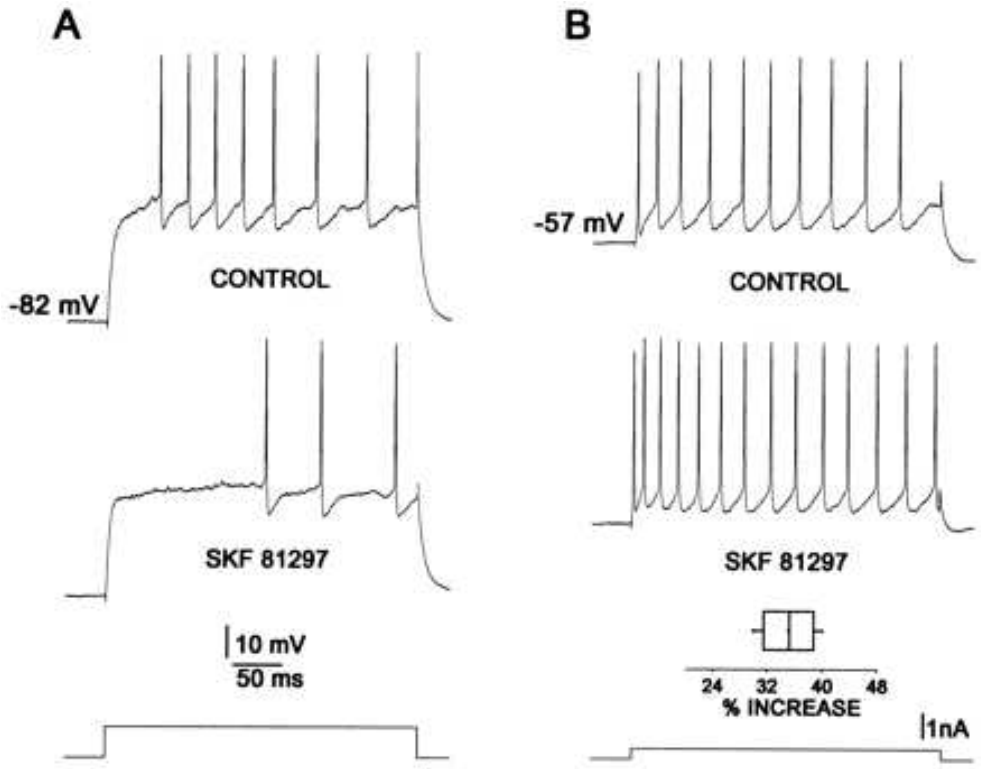

Figure 1.3: Effect of dopamine D1 receptor stimulation depends on the state of the membrane potential in the target cell. In the presence of D1 agonist SKF 81297, firing is A) reduced for cells that are in their "down-state" (resting membrane potential V_m= -82 mV), but B) increased for cells in their "up-state" (V_m = -57mV). From Hernandez-Lopez et al. (1997)

firing (raising the threshold for action and therefore less Go responding) but also enhance DA bursting (thereby enhancing learning of positive outcomes of Go responses). Consistent with this, D2 blockade during acquisition of an appetitive response actually increased behavioral output in a subsequent test conducted after drug washout (Horvitz, 2001) — that is, the drug had a NoGo performance effect during acquisition, but resulted in enhanced Go learning that was only evident once this performance effect wore off.

These dual (pre and postsynaptic) effects of D2 agents are supported by data on the dose and time-course properties of haloperidol (D2 antagonist) administration. At single low doses, postsynaptic receptors are not significantly bound, whereas the autoreceptors are far more sensitive (Richfield et al., 1989; Bannon et al., 1980), and increases in spike-dependent DA bursting in the BG are observed (while leaving DA levels relatively unaffected in PFC) (Moghaddam & Bunney, 1990). Similarly, other D2 antagonists also have much greater selectivity for presynaptic receptors at low doses and only activate postsynaptic receptors at higher doses (Schoemaker et al., 1997). With chronic administration at higher doses, catalepsy and Parkinsonism is induced, which is thought to reflect postsynaptic D2 receptor blockade (Sanberg, 1980). The delay of cataleptic onset is thought to arise from the initial increase in DA bursting caused by the drug's presynaptic effects, which oppose the postsynaptic antagonism effects (Phillips et al., 2001; Garris et al., 2003). Thus, haloperidol administration in healthy participants should have a markedly different cognitive profile (increased DA phasic bursting and increased NoGo firing) in comparison with PD (low DA in both phasic bursting and tonic firing, favoring NoGo overall).

1.1.2 D1/D2 Modulation of Synaptic Plasticity and Learning

Phasic bursts of dopamine may also have long-term effects that are not directly observable during the period in which it is released. A common notion is that dopamine modulates synaptic plasticity, so that synaptic weights change more during DA firing. Because "Hebbian" learning between two neurons critically depends on the firing rates of the neurons in question (Hebb, 1949), the fact that DA modulates neuronal excitability suggests that it also affects Hebbian learning. DA enhances learning for some cells (i.e., those whose activity is amplified), while blocking it for others. In this sense DA may act as a "teaching signal" that indicates when stimuli are particularly salient or rewarding (see section 1.2.2), and that the animal should learn a lot about them so that they are appropriately responded to in the future (Schultz, 2002).

Aside from activity-dependent learning, DA also directly modulates synaptic plasticity. As mentioned above, the assignment of dopaminergic receptors to the two general classes is determined by whether they cause a G-protein mediated increase (D1) or decrease (D2) in the level of adenylyl cyclase. Because adenylyl cyclase is known to be a critical part of a biochemical cascade that leads to long term potentiation (LTP), dopamine plays an important role in LTP. By stimulating adenylyl cyclase, D1 activity has the end effect of phosphorylating DARPP-32 (which is known to lead to LTP), whereas D2 activity restricts LTP (Nishi et al., 1997). Accordingly, LTP is blocked by D1 antagonists and enhanced by D2 antagonists (for a review, see Centonze et al., 2001). Phasic application of DA in prefrontal cortex facilitates LTP (Blond et al., 2002), and the administration of D1 agonists or activation of adenylyl cyclase enhances LTP at hippocampal-prefrontal synapses (Gurden et al., 1999; Gurden et al., 2000). Finally, these plasticity effects have been demonstrated to affect learning and are implicated in flexibly modifying cortical representations (figure 1.4, Bao et al., 2001; Bao et al., 2003).

1.1.2.1 Systems Level Effects of BG DA: Summary

Table 1.1 provides a concise summary of the core mechanisms, in terms of D1 and D2 modulatory effects on the Go/NoGo pathways and DA bursting:

- D1 modulates the Go pathway, and D2 modulates NoGo, but in opposite directions. Thus, increases in DA overall tend to excite Go and inhibit NoGo. Decreases in DA have the opposite effect, exciting NoGo.

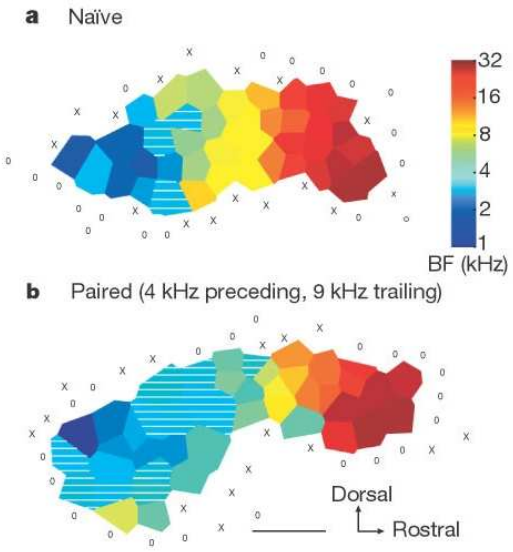

Figure 1.4: Receptive fields in the rat auditory cortex for different acoustic frequencies. External dopamine stimulation immediately following an acoustic tone results in enhanced auditory cortex receptive fields for the frequency of the tone. a) Control rats naturally have smaller receptive fields for 4 kHz (light blue) then for 9 kHz (yellow), b) Repeated stimulation of dopaminergic VTA cells after presenting a 4 kHz tone and before a 9 kHz tone has the effect of enlarging the 4 kHz receptive field and reducing the size of the 9 kHz receptive field. Adapted from Bao et al. (2001)

- Besides having a performance effect on Go/NoGo behavior, phasic changes in DA that follow good and bad responses can drive Go and NoGo learning, respectively, by modulating neural activity in Go/NoGo pathways and leading to synaptic modification. This makes the corresponding Go or NoGo representation more likely to get activated in similar subsequent experiences, and therefore supports learning.

- Presynaptic D2 autoreceptors control the level of phasic DA release during bursting. Stimulation of these autoreceptors reduces DA release, whereas their blockade leads to enhanced bursts.

1.1.3 Cellular Mechanisms of DA in the PFC: Robust Maintenance

Because the PFC is critically involved in working memory and cognitive control (Fuster, 2001; Goldman-Rakic, 1995), DA may modulate activity in the PFC that serves these functions. D1-like receptors, which are 20-fold more prevalent in PFC than D2-like receptors (Lidow et al., 1991), are particularly important for supporting PFC maintenance of neural activity (Dreher & Burnod, 2002; Durstewitz & Seamans, 2002; Durstewitz et al., 1999; Goldman-Rakic, 1996; Williams & Goldman-Rakic, 1995). In primates, injection of DA directly in PFC results in more persistent activity associated with working memory, and D1 (but not D2) antagonists reduce this activity (Sawaguchi, 2001; Wang et al., 2004).

Dopamine in the PFC is thought to increase the signal-to-noise ratio, much in the same way as described above in the BG (Durstewitz & Seamans, 2002; Cohen et al., 2002; Lewis & O'Donnell, 2000; Servan-Schreiber et al., 1990). That is, D1 stimulation results in increased excitability in response to afferent input, with a concomitant suppression of spontaneous activity (i.e., noise). Without sufficient dopamine in the PFC, task-related information is more susceptible to interference and is therefore less stable (Durstewitz et al., 2000a). On the other hand, too much dopamine may lead to overly focused representations because only a very few cells with the strongest afferent drive would get enhancement of activity, while the remainder are suppressed. Thus there is an optimal range of dopamine levels in the PFC that would support robust maintenance of top-down task-relevant information in cognitive tasks involving working memory or executive control (Goldman-Rakic, 1996; Camperi & Manias, 2003). As such, both deficient and excessive DA in the PFC lead to working memory deficits (Brozoski et al., 1979; Sawaguchi & Goldman-Rakic, 1991; Arnsten et al., 1994).

1.2 What Causes Dopamine Release? Implications for Addiction and Motivation

Clearly under natural conditions we do not have scientists injecting us with different classes of dopamine agonists or antagonists. The approach of linking cellular mechanisms of DA to behavior is only ecologically valid if we know which types of behavioral situations, mediated by which brain regions, produce increases and decreases in DA release.

1.2.1 Brain Areas Controlling DA Release

The largest dopamine system is the mesencephalic dopamine system, consisting of dopaminergic cells in the midbrain. This system can be further divided into the **mesolimbic** and **nigrostriatal** systems based on their efferent projections (Joel & Weiner, 2000). The mesolimbic dopaminergic cells are located in the ventral tegmental area (VTA) and project to the limbic striatum (including the nucleus accumbens) and frontal cortex. The nigrostriatal cells are located in the substantia nigra pars compacta (SNc) and project to the dorsal striatum.

The cells in these regions fire at intrinsic baseline levels, so that dopamine is continuously being released in all target regions. This intrinsic firing is associated with tonic dopamine release, is relatively unaffected by synaptic input, and stems from a spontaneous depolarizing pacemaker membrane potential (Grace & Onn, 1989). In addition to this "regular" mode of DA cell firing, there are also "irregular" and "bursting" modes (Hyland et al., 2002). The bursting mode is evoked by synaptic input and causes phasic increases in DA release

that may be important for learning. Interactions among limbic brain regions cause or prevent bursting: both the VTA and SNc receive afferent projections from the limbic striatum (Joel & Weiner, 2000), which in turn receives projections from amygdala, orbitofrontal cortex, and hippocampus (Groenewegen et al., 1999).

1.2.2 DA and Reward

Behavioral events related to reward are well known to induce phasic changes in DA release (Schultz, 2002; Volkow et al., 1999). While non-rewarding events such as novel, salient and even aversive stimuli may also cause DA bursting (Horvitz, 2000), converging evidence from multiple methods and in several species demonstrate that DA is most critically involved in reward signaling. However, it is now clear that phasic changes in DA are not caused by rewards themselves, but are rather elicited if **expectations** of reward are inconsistent with actual reward delivery (Schultz, 1998).

Recordings from the firing of monkey dopaminergic cells in different stages of classical conditioning paradigms have been particularly enlightening. In these studies, a stimulus (e.g., a light) precedes reward delivery, and the animal learns to associate the light with reward. Results consistently demonstrate that before conditioning, the DA neuron fires phasically just after presentation of reward (figure 1.5, Schultz et al., 1993). However, once the monkey has made a reliable association between the light and reward, the cell no longer fires during reward. Instead, it fires at the onset of the light, as if the conditioned stimulus itself now acts as the unexpected reward. If the reward is expected but not delivered, a phasic **decrease** in baseline firing is observed precisely at the point in time in which the reward was expected to occur (Satoh et al., 2003; Hollerman & Schultz, 1998).

Taken together, these observations support the notion that short term changes in dopamine firing signal violations of predictions of reward. Increases in firing indicate rewards that are delivered but unexpected, whereas decreases in firing indicate rewards that were expected but not delivered. The magnitude and reliability of these phasic changes in DA varies with the degree to which a reward is expected (Satoh et al., 2003). The more unexpected (expected) a reward is, the greater the increase (decrease) in DA firing when it is (not) received.

Phasic bursts and dips of dopamine are likely to have functional consequences. Dopamine may modulate synaptic plasticity during unexpected rewards so that animals can quickly learn what they did that got them the reward. In the BG, bursts of DA transiently activate the direct pathway, enhancing "Go" signals for the action just executed. Conversely, phasic dips in DA when rewards were expected but not received may support unlearning of the selected action, to favor the learning of some other action. This is because the indirect pathway is more activated during DA dips (see section 1.1.1), and a "NoGo" association may be learned to override the non-rewarding response.

Given the complexity and graded nature of the nervous system, it should not be surprising that some stimuli unrelated to reward also phasically activate dopamine cells. DA cells respond to novelty, even for stimuli that are not at all rewarding (Horvitz, 2000), suggesting that they might provide an alerting signal to events that should be given attention because they **might** be important. Similarly, they also respond to stimuli that are salient and attention-inducing, with bursting characteristics that match those elicited by reward (Schultz, 1999; Ljungberg et al., 1992). Finally, phasic dopamine firing is also elicited by stimuli that do not predict reward but merely physically resemble other stimuli that do. In all of these cases, DA bursts are followed by short term decreases in baseline responding (Kakade & Dayan, 2002; Schultz, 1998). It is possible that a fast signal provides general properties of reward-like information through the basolateral amygdala and limbic striatum to dopaminergic areas of the midbrain, causing DA bursts. Once the stimulus is cortically processed and its details indicate that it is in fact not a reward predictor, a second wave of amygdala-limbic striatum activity may cause a phasic dip in DA because the reward was initially expected and now it is not. This argument is similar to the direct and indirect fear pathways to the amygdala advocated by Ledoux (Armony et al., 1997; LeDoux, 1992, 1998).

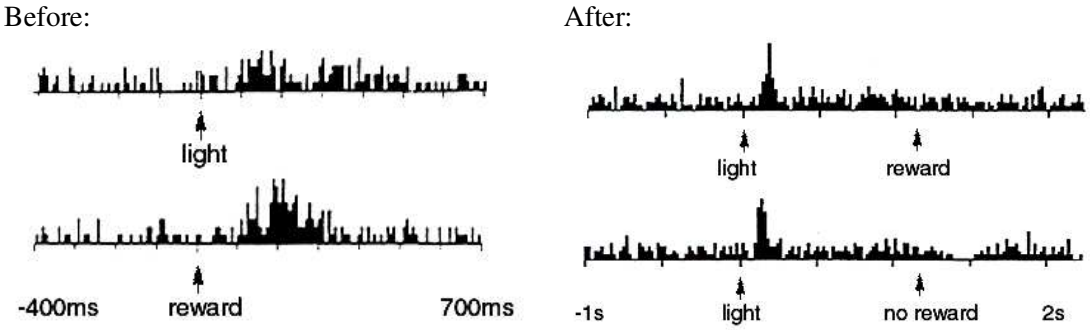

Figure 1.5: Phasic bursting of a dopaminergic cell in the ventral tegmental area (VTA) before and after a monkey was conditioned to predict reward following the presentation of a light stimulus. Before conditioning, some phasic bursting is observed when the light is presented (possibly due to its novelty), but substantially more bursting is observed just after the unexpected reward. After conditioning, the onset of the light stimulus evokes a precise burst (as it is associated with reward but is itself unexpected), but no bursting occurs during the presentation of the now expected reward. Finally, if the reward is not delivered following the light, a **decrease** in baseline firing occurs at the time the reward was predicted, indicating a violation of reward expectation. Adapted from Schultz, Apicella & Ljungberg (1993).

1.2.2.1 Dynamic Range in DA is Needed to Distinguish Between Reward Values

In the above discussion it was assumed that responses are either rewarded or not. However, a more realistic scenario is that certain behaviors are sometimes rewarding but their consequences are not entirely predictable. The ability to implicitly discriminate between responses that are rewarding 80% of the time versus those that are only rewarding 50% of the time may depend on a healthy **dynamic range** of DA release in the basal ganglia. Bursts of DA are essential for learning Go associations (by inducing Hebbian learning in the direct pathway) and dips in DA enable learning of NoGo associations (by releasing the indirect pathway from the suppressive influence of DA). If a burst is elicited 80% of the time, a stronger Go association will be developed and the animal will learn to execute the associated response.

For this system to work, the DA signal must be able to both increase and decrease substantially from baseline levels during rewarding and non-rewarding events (Frank, in press). DA depletion of the BG (as in Parkinson's disease) may prevent this from occurring because DA levels are always low and cannot increase enough to learn appropriate Go responses. Conversely, an excess of DA in the BG would tonically suppress the indirect pathway, causing learning of Go responses for both 50% and 80% reward probabilities. Thus, to learn changing reinforcement values of responses, a large dynamic range of DA release may be critical (Frank, in press). This point will be evident when discussing effects of DA medication, and DA depletion in animals and Parkinson's patients.

1.2.3 DA and Addiction

Because of its ostensible relationship with the reward system, dopamine has been implicated as a neurobiological basis for addiction. Several recreational drugs are known to increase DA release. The effects of cocaine have been particularly well studied, as it blocks the reuptake of DA into the cell, leaving extracellular DA levels far greater than baseline (Giros et al., 1996). Opiates such as heroin decrease the release of GABA, which normally inhibits the firing of dopaminergic neurons, thereby disinhibiting DA release. Thus, a prevailing assumption is that recreational drugs tap into the brain's natural reward system by artificially forcing dopamine levels to match those stimulated by actual rewards (e.g., Wise, 2002; Volkow et al., 1999). Converging evidence of this hypothesis comes from experiments in which rats learn to press a lever that applies electrical stimulation to their own substantia nigra, thereby releasing dopamine. No food or liquid reinforcement is provided in these studies; the dopaminergic release seems to be rewarding enough in and of itself to drive the animal to continue pressing the lever. Further, in support of dopamine-mediated reinforcement learning, corticostriatal efficacy is potentiated during intracranial self stimulation, and this effect is blocked by DA antagonists (Reynolds et al., 2001).

Studying the cognitive processes of abusers who are on and off drugs may provide valuable insight into the role of DA in these processes. Chronic cocaine use decreases the amount of dopaminergic cell activity, as well as the number of D2 receptors in drug abusers (Volkow et al., 2003). These effects may contribute to the positive feedback loop associated with addiction, in which the more a drug is abused, the more of it is needed to produce the same high. At the same time, the dynamic range of endogenous DA is reduced in addicts, so they are less responsive to natural rewards and are therefore less likely to pursue them. Even worse, increased striatal synaptic plasticity during cocaine elicited bursts of DA may cause abusers to ingrain their destructive behaviors into "habits" (Gerdeman et al., 2003), further damaging the DA system. For an explicit theoretical perspective that relates cognitive impairments in abusers to those in ADHD, please see section 1.6.

1.2.4 DA and Motivation

If DA is involved with reward signaling, it should also modulate motivational processes that drive animals to pursue rewards. Indeed, the implication of DA in motivation has been so extensively studied that a full survey of the literature is beyond the scope of this review. Nevertheless, a short summary is in order so that the role

of DA in the BG can be made clear. The current framework holds that DA modulates BG gating of multiple cortical areas, including those maintaining long term goals associated with motivation.

Although addictive drugs stimulate the DA system and cause reward-like sensations, DA itself is neither necessary nor sufficient for mediating pleasureful aspects of reward. While observations that DA blockade reduces operant lever pressing for food were initially linked with the idea that DA withdrawal takes the "goodness" out of the food reward (Wise, 1985a), this hypothesis has been strongly challenged. First, striatal DA is released during events that could hardly be considered pleasureful, including the encoding of new associations of aversive, and possibly even neutral, stimuli (Horvitz, 2000; Young et al., 1998). Second, the hedonistic account of DA does not fit with a variety of other data showing that dopamine-depleted animals continue to consume food when it is presented to them. Based on this data, Berridge and Robinson (1998) conclude that DA does not mediate hedonistic aspects of reward, but may instead mediate the "wanting" of food when it is not immediately available. Others have further teased apart the aspects of wanting that seem to be affected by DA depletion: primary motivation and the appetite to consume is left intact and it is really the willingness to work for food that is deficient (Salamone & Correa, 2002). This distinction is primarily based on studies in which DA-depleted rats take the easy route to obtain small amount of food rather than climb over a barrier to obtain greater amounts of food (or more preferred food). Along the same lines, DA blockade causes rats to choose small immediate rewards instead of large but delayed rewards (Cardinal et al., 2000; Wade et al., 2000). These observations may explain why chronic cocaine users, who have downregulated DA systems, often choose immediate gratification over the pursuit of long term reward (Coffey et al., 2003). Conversely, acute increases in DA induced by amphetamine tends to reduce impulsive behavior, both in rats and healthy human participants (Wade et al., 2000; de Wit et al., 2002).

There seems to be an abundance of evidence that DA dysfunction impairs motivational processes, but one should not conclude that DA implements motivation itself. Although DA effects are undoubtedly complex, motivational states are unlikely to be implemented by a single neurotransmitter. Instead, activation states of entire brain regions—including, but not limited to, orbitofrontal cortex (OFC)—affect motivation by maintaining reward/goal information over time (Hikosaka & Watanabe, 2000). Dopamine may simply modulate function in OFC and in ventral striatal areas that are interconnected with it (for a more explicit formulation of this theory, see the next section).

1.3 Relating Basal Ganglia Roles in Motor Control and Cognitive Function

1.3.1 Behavioral Effects of BG DA on Motor Systems

Perhaps the most studied effect of dopamine on human subjects is its involvement in motor control. Parkinson's Disease (PD) is a progressive neurodegenerative disease that selectively damages the dopaminergic cells of the substantia nigra pars compacta (SNc), which target striatal medium spiny neurons of the basal ganglia. As described in detail in section 1.1, the BG is involved in motor control by modulating thalamic output to premotor cortex (e.g., Wickens, 1997). The most obvious behavioral change associated with PD is a loss of motor control characterized by muscular rigidity, slowness of movements, and tremor. Note that PD patients do not have damage to motor neurons and can in fact perform movements quite smoothly under some circumstances, such as externally driven motor commands. Instead, these patients have difficulty selecting among various competing motor actions and executing the most appropriate one.

It is often suggested that depleted BG dopamine in PD leads to an imbalance of direct/indirect pathways (Albin et al., 1989), such that PD is associated with too much "NoGo" and not enough "Go", leading to slowness of movements or **bradykinesia**. In essence, depleted DA in the BG may result in raising the threshold for facilitating a motor program while continuing to suppress competing actions (Mink, 1996; Wichmann & DeLong, 2003). The observation that treatment with DA agonists and L-Dopa sometimes lead to jerking movements, or **dyskinesia** (McAuley, 2003) is consistent with this hypothesis by shifting the balance the other way and making the threshold for motor execution too low, rather than too high (Gerfen, 2003).

The effects of BG DA in the motor circuit are also likely involved in cognitive procedural learning. These are discussed extensively in chapters 2 and 4 as they are implicated as source of procedural learning deficits in Parkinson's disease.

1.3.2 Making the Leap: The Role of BG DA in Cognition is Similar to that in Motor Systems

1.3.2.1 Working Memory

Based on the general suggestions of basal ganglia involvement in prefrontal circuits made by Alexander and colleagues, we developed a computational model that explicitly formulated the role of the BG in working memory (Frank et al., 2001). We suggested that just as the BG facilitates motor command execution in premotor cortex by disinhibiting or "releasing the brakes", it may also facilitate the updating of working memory in prefrontal cortex. For task-relevant stimuli that are suitable for PFC maintenance, the BG direct pathway may activate a "Go" signal, thereby disinhibiting the thalamus and gating the updating of PFC (see figure 1.1). Task-irrelevant stimuli might still transiently activate PFC cells, but due to mostly NoGo BG output, they would not be robustly maintained.

Follow-up work demonstrated that phasic changes in DA, via its reward signaling and modulation of synaptic plasticity, could modulate the learning of task-relevant stimuli (Frank, in press). In particular, DA modulated the relative balance of direct/indirect pathways and supported Go/NoGo procedural learning, providing a parsimonious explanation for cognitive deficits in Parkinson's disease (see chapter 2). According to this model, DA bursts in the BG should preferentially activate (and increase the signal-to-noise ratio of) cells in the direct pathway via D1 receptors, while suppressing cells in the indirect pathway via D2 receptors. Increases in DA during reinforcement may also boost the updating of working memory, by biasing the direct pathway to win the competition for BG output and facilitating recurrent activity in PFC. When reinforcement is not received, a phasic DA dip would allow the BG to learn **not** to update task-irrelevant information. Once information is updated into PFC, the role of DA in that structure may be quite different, supporting the robust maintenance of information over time and in the face of interfering stimuli, as briefly reviewed in section 1.1.3.

With the above model in mind, consider the effect of dopaminergic dysfunction in the BG or PFC. A lack of DA in the BG would lead to too little updating of relevant information into PFC, just as it leads to too little execution of motor commands. Conversely, too much DA in the BG would lead to excessive updating of PFC, just as it leads to L-Dopa induced motor tics and dyskinesia in Parkinson's disease. Finally, a suboptimal level of DA in the PFC would lead to insufficient maintenance of task-relevant information. Any of these DA dysfunctions would lead to "frontal-like" cognitive deficits.

Other models also emphasize the need for a dynamic gating mechanism, but contend that this is implemented by phasic DA bursts directly in PFC (Braver et al., 1999; Braver & Cohen, 2000; Cohen et al., 2002). Specifically, these authors propose that the reward signaling of DA bursts transiently increases the efficacy of PFC afferents from posterior cortex, leading to updating. This is consistent with the current framework in that bursts of DA induce PFC updating, but the proposed mechanism is not biologically plausible (Dreher & Burnod, 2002). First, if DA bursts in PFC support updating, one might expect post-synaptic effects of D1 stimulation to be short-lasting. However, D1 stimulation effects in PFC are delayed and temporally extended (Seamans et al., 2001), making them unlikely to support fast updating of activity (Seamans & Yang, 2004; Dreher & Burnod, 2002). To reconcile this data, Cohen, Braver, and Brown (2002) proposed that phasic DA exerts its effects in PFC via D2 stimulation. Recent evidence shows that D2 stimulation indeed affects phasic activity in PFC (Wang et al., 2004). However, these effects were specifically **at the end** of the delay period, during initiation of motor response, which is not consistent with the updating of working memory. Further, because DA release in PFC is diffuse, a role in updating would imply that each burst induces updating of all PFC subregions simultaneously. This has poor computational properties, because it does not allow for **selective** updating of some information, with concurrent maintenance of existing information already stored in PFC. Selective updating is necessary in many aspects of executive function, in which higher level goals should be maintained while lower

level subgoals can be continuously updated (Frank et al., 2001; Cohen et al., 2002).

This problem is avoided if the BG performs the gating function. Because parallel sub-loops exist between BG and PFC, a Go signal in one part of the BG may enable updating of information only in a selective region of PFC (for more discussion of this issue, see Frank et al., 2001; O'Reilly & Frank, in press; Frank, in press). DA simply facilitates the execution of the Go signal by suppressing NoGo activity (see section 1.1.1). Further, rapid updating is plausible in this scenario because the Go/NoGo balance is largely mediated by post-synaptic effects of D2 stimulation, which unlike those of D1, are short-lasting (Seamans et al., 2001).

1.3.2.2 Motivation and Impulsiveness

As described in detail in section 1.2.4, stimulation and blockade of DA systems has a profound effect on an organism's tendency to compare incentive values of their actions. Whereas DA blockade is associated with low motivation and increased impulsiveness, DA stimulation has the opposite effect. Phasic changes in DA signal low-level aspects of reward, and resulting effects on motivation and impulsiveness are related to processing in ventral striatum and orbitofrontal cortex (Depue & Collins, 1999).

The current framework holds that rewards and their predictors cause phasic bursts of DA in ventral striatal areas which, beside facilitating motor responses in motor cortex, may also facilitate the updating and subsequent maintenance of reward information in OFC. Just as DA lowers the threshold for dorsal striatal gating of motor commands, so it may lower the threshold for ventral striatal gating of reward information in OFC. Maintained reward information may then guide goal-directed behavior via interactions with other areas of prefrontal cortex and by directly controlling striatal motor systems (Wallis & Miller, 2003), over and above its potential role in modulating DA reward-processing systems.

Several lines of evidence support the hypothesis that ventral striatum gates reward information to be maintained in OFC. First, well established cortico-striatal circuits link these regions (Alexander et al., 1986; Haber et al., 1995). Second, the BG is thought to gate information in frontal cortex (Mink, 1996; Chevalier & Deniau, 1990; Frank et al., 2001), and the content of that information depends on what is normally represented in the cortical region of interest (Frank, in press). Via afferent projections from the basal lateral amygdala (Cavada et al., 2000), the OFC represents reward and incentive value information which it maintains in working memory (Schoenbaum & Setlow, 2001). Finally, indirect evidence of the functional interactions between striatum and OFC comes from cocaine or methamphetamine users. Chronic use of either drug diminishes DA processes in the striatum (Volkow, Fowler, & Wang, 2003), which according to the current framework, should result in reduced frontal gating. Indeed, reduced striatal DA in abusers has been linked with reduced metabolic activity in OFC (Volkow et al., 2001a) and cingulate cortex (Kaufman et al., 2003). Notably, these abusers are impaired in their ability to maximize long term rewards, likely due to altered OFC function (Grant et al., 2000; Bolla et al., 2003).

That the OFC monitors reward values of behaviors is consistent with both animal and human literature. In animals, the OFC is critical for learning stimulus-reinforcement contingencies, particularly when these contingencies change over time, as in reversal learning (e.g., Rolls, 1996). In humans, it is generally thought that the OFC is important for long-term decision making and reducing impulsiveness (Bechara et al., 1998; Mobini et al., 2002). These theories implicitly assume that the OFC must somehow reduce the motor system's habitual tendency to execute prepotent or short-sighted actions. An intriguing possibility is that the maintenance of long-term goals in OFC suppresses the ventral striatum from impulsively facilitating the execution of prepotent responses. One way it may accomplish this function is by controlling striatal DA release. It has been shown that PFC (and by extension, OFC) stimulation inhibits nucleus accumbens DA release (Jackson et al., 2001)[3] , suggesting that endogenous OFC firing effectively raises the BG gating threshold and may therefore have the desired effect on motor output.

[3] Although others report increased nucleus accumbens DA release resulting from PFC stimulation (You et al., 1998), these studies involved stimulation at much higher frequencies than occur naturally. It is also possible that PFC stimulation increases tonic DA levels which then act presynaptically on autoreceptors to inhibit phasic release (e.g., Grace, 1995).

Given the above, we can speculate that long-term reward information maintained in OFC inhibits further phasic bursts of striatal DA so that short-sighted responses are overridden. While prepotent affective associations tend to elicit motor responses, OFC activity can control these emotional reactions to some extent. In animal models, affective associations induced by external amygdala stimulation indeed increase accumbens DA release (and therefore motor output), but only if PFC transmission is blocked (Jackson & Moghaddam, 2001). Similarly, injecting DA agonists in rat medial PFC—homologous to primate OFC—attenuates the accumbens DA response to stress (Stevenson & Gratton, 2003). That OFC encodes relative reward preference (Tremblay & Schultz, 1999) further supports this idea, as anticipation of greater rewards should have a suppressive effect on behaviors that would lead to less valuable rewards.

Note that this hypothesis does not imply that OFC is globally inhibitory. Indeed, the suppression of prepotent responses must be accompanied by a facilitation of responses that will result in larger, longer term rewards. While external PFC stimulation seems to inhibit striatal DA release (Jackson et al., 2001), it is also possible this reflects the fact that the "average" PFC representation has an overall inhibitory effect but that specific representations can actually have an excitatory effect. If the **content** of information is taken into account, goal-directed PFC representations may actually facilitate DA release, as the nucleus accumbens learns that they are rewarding. Moreover, the suppressive account of OFC function advocated above may apply to lateral areas of OFC, whereas medial regions are more involved in learning new stimulus-reward associations (Iversen & Mishkin, 1970; Elliott et al., 2000). Consistent with this distinction, stimuli with low reward value activate lateral areas of OFC while those with high value activate more medial regions (O'Doherty et al., 2001). A tentative OFC division of labor is that medial OFC maintains the most rewarding stimuli/responses to guide behavior over the course of a task and directly biases response-related activity in other frontal (.e.g, premotor) areas. Concurrently, lateral OFC may support NoGo firing in the BG to suppress less rewarding striatal habits.

1.3.3 Summary

This section summarized the primary hypothesis of the dissertation, which was motivated by both biological characteristics of BG-PFC circuitry as well as computational constraints on the functional components necessary for a working cognitive system. In brief, the BG is thought to provide a complementary role to the PFC in cognition, by controlling when representations in PFC are updated. In addition to its implications in working memory, this hypothesis is also relevant in motivational processes that depend on orbitofrontal areas. Dopamine dynamically modulates the BG threshold for updating frontal representations, while tonically enhancing maintenance properties of PFC itself. In the remainder of the chapter I will present converging evidence for this hypothesis, before moving on to computational and empirical tests of its plausibility.

1.4 Concurrent Development and Decline of DA System and Cognitive Control

1.4.1 Phylogeny

It has been hypothesized that the origin of human intelligence and executive control is related to the evolutionary expansion of dopaminergic systems, in terms of more widespread projection of dopamine fibers in different regions and within different cortical layers of the brain (Previc, 1999). A series of arguments demonstrated that changes in size and shape of the brain typically attributed to giving rise to human intelligence may instead be byproducts of it. Rather, Previc argues, the emergence of chase-hunting in sub-Saharan Africa may have led to a greater consumption of protein-rich meat, adding tyrosine to the diet, which was easily converted to dopamine due to the increased physical activity.

1.4.2 Ontogeny

Both the PFC and the DA system continue to develop into adolescence. Using diffusion tensor MRI, it was demonstrated that there was less myelination of PFC white matter in 8-12 year old children, relative to adults (Klingberg et al., 1999). Frontal development occurs with concurrent improvement in executive tasks involving cognitive control, behavioral inhibition, and working memory (Diamond & Goldman-Rakic, 1986, 1989; Diamond, 1990; Diamond & Taylor, 1996; Hale et al., 1997; Carver et al., 2001).

In monkeys, the dopaminergic system continues to develop its axonal innervation of the PFC until ages 2-3, which roughly corresponds to ages 10-12 in humans (Rosenberg & Lewis, 1995). In an analog of Piaget's A-not-B task in macaques, performance improves throughout the first 6 months of life, concurrently with rising DA levels and DA receptor gene expression (Goldman-Rakic & Brown, 1982; Lidow et al., 1991)

Other evidence of the involvement of DA in the development of cognitive processes comes from children with phenylketonuria (PKU), a genetic mutation that causes decreased levels of the dopamine-precursor tyrosine. The normal drastic improvement in the A-not-B task seen in the first year of life, is not observed in PKU children (Tam et al., 1990; Thierry et al., 1977)

1.4.3 DA, Aging and Cognitive Control

After development, dopamine activity declines with age. PET studies have shown that the number of D2 receptors decreases by approximately 8% per decade after age 20 (Volkow et al., 1996b), and DA transporter levels also decrease with age (Volkow et al., 1996a). The decline of dopaminergic activity is correlated with decline in executive control in both Parkinson patients and with age in healthy individuals (Rinne et al., 2000; Volkow et al., 1998).

Of course, correlational measures of the development and decline of dopaminergic systems with that of cognitive function are not in themselves convincing, as dopaminergic integrity is likely to correlate with integrity of many other neural systems as well. It is therefore necessary to examine the functional implications of DA systems while "controlling for" these other systems. This can be accomplished (with some interpretative cautions) with lesion studies in animals, neuropsychology studies in humans with deficient DA systems, functional imaging with healthy participants during tasks that involve DA, and acute pharmacological manipulations of DA.

1.5 Dopaminergic Lesion Effects on Cognitive Function in Experimental Animals

Dopamine may have different functions (or at least different behavioral effects) in distinct brain regions. Therefore, systemic depletion of DA associated with antagonists or lesions of dopaminergic neurons that project to multiple regions may confound the interpretation of associated cognitive deficits. One way of examining the effects of DA in a brain region of interest is to selectively lesion the dopaminergic terminals in that region, by infusing the toxin 6-hydroxydopamine (6-OHDA), and leaving dopaminergic processes intact elsewhere. This technique has revealed dissociations between effects of DA lesions in different regions. DA depletion in nucleus accumbens led to locomotor hypoactivity, whereas that in cortex led to hyperactivity (Jones & Robbins, 1992). In the context of the present framework, a lack of DA in the NAc should indeed result in NoGo BG output to motor cortex. As for cortical DA depletion, it is now known that this often results in a reactive **increase** in striatal DA (Nieoullon, 2002), which may explain the opposite effect.

The effects of DA depletion on cognitive processes are reviewed next, for both rodents and primates. It should be kept in mind, however, that there are two disadvantages to the 6-OHDA technique. First, the lesions are not 100% complete and some neural regeneration is possible (Collins et al., 2000), so that null effects found with these lesions are not at all conclusive. Second, as mentioned above, DA depletion of one brain region may be associated with compensatory increases of DA in other areas, which may themselves cause behavioral changes. Indeed, a current theory of the neurobiology of schizophrenia is that the "negative"

symptoms are caused by dopamine depletion of the prefrontal cortex, whereas the "positive symptoms" are caused by a reactive increase of dopamine subcortically (Weinberger, 1987).

1.5.1 Rodents

6-OHDA lesions of the dorsolateral (motor segment) striatum of rats produces slowness of movements, similar to the bradykinesia seen in Parkinson's patients (Amalric et al., 1995). The effects of DA lesions in other parts of the striatum are complex, supporting a more cognitive role. In particular, striatal DA may be important for distinguishing between multiple responses and flexibly learning changes in rewarding behaviors over time (see section 1.2.2.1). To learn a new behavioral response, sufficient DA must be available to bias the direct/Go cells to facilitate its execution. Consistent with this hypothesis, depleted striatal DA impaired switching from one situation to another with increased reaction times and perseverative behavior in a five-choice serial reaction time task (Baunez & Robbins, 1999).

In addition to distinguishing between multiple previously-rewarded responses, striatal DA may be involved in optimally determining **when** to initiate a single response. Rats were trained to depress a lever and release it when a light stimulus was presented, which occurred after variable delays from 500 to 1250 ms. In intact rats, the reaction time to the light stimulus decreased with increasing delay. In other words, as time elapsed, the animals could use temporal information in order to prepare motor responses. This decreasing reaction time profile was not observed in rats with DA-depleted striatum, despite their normal RT's for simple motor responses (Amalric et al., 1995). These results were interpreted to implicate DA in the BG as being important for timing. However, it is possible that timing information comes from elsewhere (e.g., the cerebellum) and that striatal DA is necessary for distinguishing among rewarding event probabilities and their associated behavioral reactions. Note that in this paradigm the light stimulus was sure to occur at some point after the tone, and the probability of it occurring at any point in time increased with increasing delay. DA levels in the BG may gradually increase as time elapses, especially if the timing is uncertain (Fiorillo et al., 2003). This may have the effect of gradually biasing the direct pathway to be more active, making the animal more responsive as time elapses.

1.5.2 Primates

Experimental work with non-human primates provides a way of examining the neural processes associated with higher level cognition that are more likely to capture the complexities underlying human thought. In particular, the neural mechanisms of working memory (WM) have been primarily investigated in primates, as they demonstrate human-like abilities to maintain multiple task requirements on-line and, unlike rodents, they have a well defined prefrontal cortex (PFC). Thus the model put forth in section 1.3.2 can be critically examined in monkeys, which suggests that (a) DA in the BG is necessary for updating working memory by allowing direct pathway to win the competition for BG output against the indirect pathway; and (b) DA in the PFC is essential for working memory maintenance.

Commonly studied is the spatial-delayed response task, in which monkeys have to recall the location of a visual stimulus presented at the beginning of a trial over a delay period during which no stimulus is present. It has been well documented that PFC integrity is necessary for good performance and that PFC cells continue to fire during the delay period, apparently encoding the to-be-recalled location (Goldman-Rakic, 1996).

The involvement of DA in these tasks has also been well studied. Increases of extracellular DA were found in monkey PFC during a WM task compared to baseline and to a non-WM task (Watanabe et al., 1997). As alluded to in section 1.1.3, there is an optimal range of DA levels in the PFC that would support robust maintenance of top-down information in WM tasks (Goldman-Rakic, 1996). As such, depletion of DA in the PFC causes deficits in spatial-delayed response tasks (Brozoski et al., 1979; Sawaguchi & Goldman-Rakic, 1991), but so does excessive DA (Arnsten et al., 1994). Notably, 6-OHDA lesions to the caudate nucleus results in severe impairments in a spatial delayed response task, similar to those observed in monkeys with full

prefrontal lesions (Collins et al., 2000), supporting a role of BG DA in modulating frontal function.

Another study examined differential effects of 6-OHDA lesions to the frontal cortex and caudate nucleus in acquiring an attentional set (Crofts et al., 2001). These authors found that DA depletion to frontal cortex caused increased distractibility to task-irrelevant items, presumably due to decreased top-down support of task-relevant cues in working memory. In contrast, DA depletion of the caudate made monkeys **less** distractible (i.e., they were less sensitive to changes in task-irrelevant cues). This can be understood in the context of the role of striatal DA in setting the threshold for updating PFC, put forth in section 1.3.2: less DA results in a higher threshold, making irrelevant information less likely to get updated and therefore less distracting. Of course, the limited nature of 6-OHDA lesions might play a role in this effect—a complete DA lesion in the striatum would be expected to raise the threshold such that even task-relevant items do not get updated. Further, while DA depletion of the BG results in less distractibility in this paradigm, it may also force the animal to be too rigid in what to pay attention to.

In this regard, caudate DA depletion causes deficits in the shifting of attentional sets, similar to those observed in monkeys with frontal damage (Dias et al., 1997; Collins et al., 2000). However, there is an interesting dissociation in attentional deficits seen after frontal and basal ganglia damage. Frontal lesions impaired shifting from one set to another, but no deficit was observed when the subject had to then make a second shift back to the originally relevant set (Dias et al., 1997). In contrast, BG DA depletion resulted in spared performance on the first shift, but monkeys were impaired at shifting **back** from the second set to the first (Collins et al., 2000). Further research is necessary to conclusively evaluate the significance of this dissociation, but a tentative speculation is that in the first shift, the PFC provides top-down support to bias the new relevant attentional set so that it can properly compete with prepotent bottom-up associations of the previous set (O'Reilly et al., 2002). In contrast, the BG may learn Go and NoGo to the different sets to decide whether they will be updated and maintained in PFC. Note that because the first shift requires learning NoGo to the originally relevant set, the second shift requires overriding this NoGo association. BG DA depletion results in hyperactivity of the indirect/NoGo pathway which may make it possible to **learn** NoGo but much more difficult to override this association and switch back to Go.

Whereas DA integrity in the BG seems to play a role in executive control and attentional set-shifting, the role of DA in the PFC is less clear for these processes. In contrast to the deficits seen after excitotoxic PFC lesions, 6-OHDA lesions in the PFC of monkeys actually resulted in **enhanced** performance on an analog of the Wisconsin Card Sorting Test (Roberts et al., 1994). Notably, **in vivo** microdialysis revealed a reactive increase in striatal DA which may have contributed to improvement in attentional set-shifting required by the task.

In summary, DA function within the BG appears to be critical for normal cognitive function. A mechanistic framework for understanding BG/DA function in animals may provide informative constraints toward understanding how this framework may apply to humans. In the next chapter, I will address how these ideas may inform theoretical perspectives for understanding complex cognitive deficits in patients with Parkinson's disease (Frank, in press).

1.6 Cognitive Impairments in ADHD

1.6.1 ADHD: Dopamine Deficiency?

In addition to Parkinson's Disease (elaborated in the next chapter), striatal dopaminergic dysfunction is also thought to form the neural basis for ADHD (e.g., Sagvolden et al., 2004; Biederman & Faraone, 2002; Solanto, 2002). However, whereas PD stems from a lack of available DA, there is no shortage of DA in ADHD. Instead, an overactive reuptake system may remove too much extracellular DA from the synapse, as both children and adults with ADHD have abnormally high densities of dopamine transporters (DAT's) (Dougherty et al., 1999; Krause et al., 2000). By analyzing the neural implications of these differences, we can explain the grossly distinct behavioral profiles observed in the two conditions. In PD, a lack of DA ensures that both tonic and phasic levels are low. But in ADHD, low levels of tonic DA may actually result in hypersensitivity

to phasic bursts (Grace, 2001). According to the current framework and supported by additional evidence described below, hypersensitivity to phasic DA bursts in ADHD causes impulsive and hyperactive behavior by transiently enhancing BG Go signals and suppressing NoGo signals.

Although it could be argued that the increased DAT expression in ADHD is an epiphenomenon, or a compensation of some other underlying disorder, genetic studies confirm that ADHD is associated with a polymorphism in the DAT1 allele (Cook et al., 1995; Gill et al., 1997; Roman et al., 2004). A positive correlation found between the expression of the DAT1 allele and scores of hyperactivity-impulsivity in ADHD (Waldman et al., 1998) suggests that this is behaviorally relevant. However, these findings only account for 3.6% of the variance in hyperactive-impulsive symptoms, so the increased DAT expression in ADHD reported above may not be exclusively genetic. Nevertheless, other genetic polymorphisms may have the same effect, such as those reported with regard to D2 and D4 receptors (e.g., Biederman & Faraone, 2002; Rowe et al., 1998; Comings et al., 1996), which could conceivably cause a reactive change in DAT expression.

Stimulants such as Ritalin (methylphenidate) have paradoxical calming effects on many ADHD patients. This is likely related to their known efficacy in blocking the DAT (Volkow et al., 1995). By blocking a percentage of DAT's, stimulants elevate tonic DA levels (Volkow et al., 2002b; Volkow et al., 2001b) and thereby reduce the potency of phasic DA bursts (Seeman & Madras, 2002). Some might find it of grave concern that in blocking DAT's, Ritalin exactly mimicks the mechanism by which cocaine produces its addictive sensations (Volkow et al., 1995; Vastag, 2001). In fact, Ritalin is just as potent as cocaine in terms of the number of transporters blocked (a typical dose given to children blocks 70% of DAT's). However, the reason for it not producing an addictive "high" is potentially related to the slow time course—unlike cocaine, which immediately increases extracellular DA levels and produces reward-like sensations, oral administration of Ritalin only gradually increases DA levels (Volkow et al., 2002a). Injecting a liquid form of Ritalin produces a potent high which cocaine addicts have likened to the feeling they get from sniffing cocaine (Volkow et al., 2002a), presumably due to a more phasic burst of DA.

1.6.1.1 ADHD: Novel BG/DA Hypothesis

While a complex disorder such as ADHD is unlikely to be a function of any single neurotransmitter, DA dysfunction of some sort—whether genetic, environmental, or a combination—is relatively undisputed. In a recent comprehensive review of the behavioral and biological bases of ADHD, the authors concluded that hypodopaminergic function in three striato-cortical loops are responsible for core deficits in DA-mediated reinforcement and extinction (Sagvolden et al., 2004). But whereas these authors claim that low levels of tonic DA are accompanied by a diminished phasic component, they cite evidence that phasic DA signals are actually **heightened** when tonic levels are too low (Grace, 2001; Solanto, 2002). This is because low levels of tonic DA are insufficient to activate autoreceptors that would normally inhibit a certain amount of phasic DA release. By raising tonic DA levels, stimulants may activate these autoreceptors and therefore reduce the system's susceptibility to phasic bursts and associated impulsiveness (Seeman & Madras, 2002).

A complementary proposal is that in ADHD low levels of tonic DA leave a greater percentage of postsynaptic DA receptors unoccupied, and therefore available for stimulation during phasic bursts. Indeed, children with ADHD were found to have increased D2 receptor availability, using Single Photon Emission Computed Tomography (SPECT) (Ilgin et al., 2001). The implication is that when endogenous bursts occur, they stimulate a greater number of receptors, sensitizing their effects on neuronal excitability.

Within the current framework, the hypothesized heightened effects of phasic DA bursts in ADHD may be particularly critical in the basal ganglia, in which they increase Go signals while suppressing NoGo signals. Behavioral symptoms may arise from lowered effective thresholds for executing motor responses and updating working memory. With too low of a threshold, impulsive behavior, thoughts, and increased distractibility may arise. Notably, the increased D2 receptor availability discussed above was found specifically in the BG (Ilgin, Senol, Gucuyener, Gokcora, Atavci, & Sener, 2001). BG cells expressing the D2 receptor, prevalent in the indirect/NoGo pathway, should then be particularly sensitive to phasic bursts of DA. Because DA is inhibitory

on D2 receptors (section 1.1.1), a D2 hypersensitivity to phasic DA bursts suggests that NoGo activity is too easily inhibited, which may make it particularly difficult to inhibit responses and pay attention to a single task. In this regard, the "dopamine hypothesis" of schizophrenia also implicates an overactive D2 system, which is often treated with D2 antagonists (Seeman, 1987). Notably, ADHD-like symptoms are observed in young at-risk relatives of schizophrenia patients (Keshavan et al., 2003).

If this hypothesis is correct, we would predict that the patient's level of BG D2 receptor availability is predictive of his/her response to stimulants. Indeed, Ilgin et al. (2001) found that Ritalin treatment was most effective for the children with the highest initial D2 receptor availabilities, which were then reduced to normal levels by the medication. Thus the current framework suggests that a high level of D2 receptor availability may result in excessive BG gating of cortical commands, because NoGo activity is overly suppressed during endogenous DA bursts. Stimulants treat this problem by raising tonic DA levels, reducing D2 receptor availability, and decreasing sensitivity to phasic bursts.

1.6.2 ADHD: Deficits in Cognitive Control

ADHD is characterized by impulsive and hyperactive behavior. These attributes are not solely based on observations made by parents and teachers — laboratory experiments have demonstrated a specific impairment in response inhibition (Barkley, 1997). The "stop signal" paradigm provides a particularly good test of inhibitory control that is dissociated from other measures of executive function (Quay, 1997), and performance is anti-correlated with clinical measures of impulsivity. In this version of a Go/NoGo task, a tone indicates when participants should press a button. On a minority of trials, a second "Stop" tone follows the first by some delay, indicating that the initiated response should be inhibited. Children with ADHD require shorter delays between the two tones in order to successfully inhibit the prepotent response, revealing an impairment in response inhibition (Logan et al., 2000).

Notably, the same stop signal deficits are observed in chronic cocaine users (Fillmore & Rush, 2002), further implicating dysfunctional dopaminergic processes as their source. But whereas the current framework holds that phasic bursts are heightened in ADHD, the opposite may be true in chronic cocaine use, due to reduced DA activity and downregulated D2 receptors (Volkow et al., 2003). Therefore while response inhibition deficits in ADHD may be caused by impulsive motor facilitation, those of cocaine users may be attributed to reduced updating of frontal cortex (see section 1.3.2.2), less maintenance, and therefore less top-down cognitive control. The distinction between these conditions is relevant because it predicts that acute stimulant treatment has opposite effects on response inhibition deficits in ADHD and drug abusers. That is, in both cases stimulants should raise tonic levels and therefore reduce the sensitivity to phasic bursts. Reduced sensitivity to these bursts may alleviate symptoms in the hypersensitive ADHD patient, but should further impair those of the "hyposensitive" drug abuser.[4] This reasoning is supported by observations that ADHD stop-signal deficits are overcome by moderate doses of Ritalin (Tannock et al., 1995), whereas those of cocaine abusers are worsened by acute doses of oral cocaine (Fillmore et al., 2002). However, the time course of stimulant effect on DA is critical: oral administration leads to a slow absorption and therefore gradual increases in tonic DA. More rapid ingestion, such as nasal or intravenous, may immediately enhance DA signals and therefore act as phasic bursts. Thus cocaine use by these methods may actually have a short-term beneficial effect on cognitive control, by enhancing DA bursts in ventral striatum and leading to more gating of frontal cortex. Recent results are consistent with this notion, as cognitive control capabilities were actually improved in cocaine users when given intravenous cocaine (Kaufman & Garavan, 2004).

Although typically considered a frontal disorder, response inhibition deficits in ADHD may be partially mediated within the BG. First, the BG is known to be specifically involved in suppressing responses (reviewed

[4] Note that by raising tonic DA levels, acute stimulant treatment should actually enhance PFC maintenance. While enhanced maintenance might be expected to improve cognitive control, this should only be the case if the **content** of PFC representations—dictated by when they were last updated—is appropriate. Reduced sensitivity to phasic DA bursts in the BG should result in less discrimination between what should (Go) and should not (NoGo) be maintained.

above). Second, DA transporters are much more abundant in BG than in cortex so the therapeutic effects of Ritalin are more likely to occur in BG (Solanto, 2002). Third, using event-related fMRI it was shown that healthy children activate areas of the BG (caudate and globus pallidus) during trials in which responses are to be inhibited, whereas children with ADHD do not (Durston et al., 2003). It has been suggested that ADHD is a selective disruption to the direct pathway of the BG, leading to interruption of cortically mediated behaviors that would normally be facilitated by BG gating (Casey et al., 2001b). These authors further implied that hypofrontality observed in ADHD may lead to a lack of top-down support of the direct pathway in the BG.

While this is generally consistent with the BG-PFC model described above, it does not consider the importance of DA. In light of the above discussion on ADHD hypersensitivity to phasic DA bursts in the BG, it is possible that hypofrontality stems from inappropriate BG gating rather than the other way around. That is, too much BG gating may cause increased distractibility and therefore less maintenance of relevant context. However, because the interactions between BG and PFC are bidirectional, it seems likely that the two possibilities are indistinguishable and may in fact co-exist. Impulsive gating would result in poor PFC maintenance over time, which in turn provides poor top-down signals from the PFC to the BG.

Dissociations have been found between BG/PFC disorders in two distinct measures of cognitive control (Casey, 2001). In the **stimulus selection** task, successful performance on a given trial requires inhibition of a previously attended stimulus attribute. Note that this task requires maintaining the previous relevant/irrelevant attributes in memory, and may involve inhibition at the stimulus level, but does not require inhibiting a motor response. The stop signal task, on the other hand, does require inhibiting a response, but does not involve working memory. Schizophrenic (SZ) children were impaired at the frontal-dependent stimulus selection task but not the stop signal task, while those with Tourrette syndrome (TS) displayed the opposite pattern of results. Interestingly, children with ADHD were impaired at both tasks, suggesting a response inhibition impairment due to BG dysfunction, and a working memory impairment, ostensibly due to hypofrontality. This suggests that ADHD is not only associated with impulsive actions (as in TS), but also impulsive **thoughts** (as in SZ). Too low of a gating threshold in the BG may cause both effects — the former being mediated by impulsive facilitation of commands in motor cortex and the latter by the impulsive updating of PFC.

Chapter 2

Using the Framework to Model Cognitive Deficits in Parkinson's Disease

[1] In cognitive neuroscience, brain regions are often characterized as if they implemented localized functions, with relatively little treatment of interactive effects at the network level. In part, this is because interactions are difficult to conceptualize and mileage has been gained from simpler theories. In some cases, however, these theoretical accounts need to be reconsidered. Some brain regions exert their effects only by modulating function in other regions and therefore do not directly implement a cognitive process. This problem is even more elusive when considering effects of neuromodulators in a single brain region, which may have indirect but substantial effects on network dynamics.

This issue applies particularly well to the effects of dopamine (DA) in the basal ganglia (BG), which are critical for many aspects of cognition (Nieoullon, 2002). Because stimulus-response (SR) tasks recruit the BG, many researchers assume that its function is to encode detailed aspects of SR mappings (e.g., Packard & Knowlton, 2002). Others advocate a subtly different modulatory role of the BG to facilitate or suppress SR-like associations that are represented in cortex (Hikosaka, 1998; Mink, 1996). This chapter explores the latter hypothesis and further suggests that DA dynamically modulates activity in an already modulatory BG, as DA levels change in response to different behavioral events. These double modulatory effects are complex and difficult to conceptualize, motivating the use of computational modeling to make them more tenable. In so doing, the model ties together a variety of seemingly unrelated cognitive deficits stemming from DA dysfunction in the BG, as in Parkinson's disease (PD).

The cognitive deficits in PD can be divided into two general classes: those that are "frontal-like" in nature, and those that reflect impairments in implicit learning. On the one hand, patients are impaired at tasks involving attentional processes or working memory (Partiot, Verin, & Dubois, 1996; Gotham, Brown, & Marsden, 1988; Dubois, Malapani, Verin, Rogelet, Deweer, & Pillon, 1994; Woodward, Bub, & Hunter, 2002; Henik, Singh, Beckley, & Rafal, 1993; Rogers, Sahakian, Hodges, Polkey, Kennard, & Robbins, 1998). Implicit learning deficits, on the other hand, do not implicate frontal processes because they generally do not involve working memory or conscious knowledge of task demands, and frontal patients do not have such deficits (Knowlton, Mangels, & Squire, 1996). Yet, PD patients are impaired at implicit sequence learning and implicit categorization (Jackson, Jackson, Harrison, Henderson, & Kennard, 1995; Ashby, Noble, Ell, Filoteo, & Waldron, 2003; Maddox & Filoteo, 2001). Similar impairments are observed in probabilistic classification, in which participants integrate over multiple trials to extract statistical regularities of the category structure (Knowlton, Squire, & Gluck, 1994; Knowlton et al., 1996). The involvement of dopamine (DA) in these tasks is not straightforward, as dopaminergic medication has both positive and negative effects on cognitive function in PD (Gotham et al., 1988; Swainson, Rogers, Sahakian, Summers, Polkey, & Robbins, 2000; Cools, Barker, Sahakian, & Robbins, 2001, 2003).

Because the neuropathology of PD involves damage to dopaminergic cells in the BG (Kish, Shannak, & Hornykiewicz, 1988), the predominant explanations for the two classes of deficits have been (a) that the damaged BG is interconnected in a functional circuit with prefrontal cortex (Alexander et al., 1986; Middleton &

[1] A version of this chapter is in press in the **Journal of Cognitive Neuroscience** (Frank, in press)

Strick, 2000b), thereby producing frontal deficits, and (b) due to damage to a "neostriatal habit learning system" (Knowlton et al., 1996; Hay, Moscovitch, & Levine, 2002). Finally, the selective cognitive impairments resulting from dopaminergic medication have been attributed to an "overdose" of dopamine in regions of the BG that are relatively spared in PD (Gotham et al., 1988; Cools et al., 2001).

To further understand the role of the BG as a functional cognitive unit, a more mechanistic explanation involving its neurobiology, and specifically the role of dopamine, is required. What is the role of DA in the BG in modulating frontal processes, and how is it involved in habit learning? This chapter accounts for Parkinson deficits by incorporating biological considerations described in the previous chapter into an explicit neural network model to test their potential role in cognitive function. In particular, two main populations of cells in the striatum respond differentially to phasic changes in DA thought to occur during error feedback. This causes the two groups of striatal cells to independently learn positive and negative reinforcement values of responses, and ultimately acts to facilitate or suppress the execution of commands in frontal cortex. Because these cortical commands may differ widely in content, damage to BG DA gives rise to seemingly unrelated deficits.

One of the network's key emergent properties is that a large dynamic range in DA release is critical for BG-dependent learning. That is, the DA signal has to both be able to increase and decrease substantially from its baseline levels in order to support discrimination between outcome values of different responses. This dynamic range is reduced in PD, accounting for cognitive deficits. The model further suggests that by tonically increasing DA levels, dopaminergic medications might restrict this dynamic range to always be at the high end of the DA spectrum, adversely affecting some aspects of cognition. For simplicity, only cognitive procedural learning tasks are modeled, but the same arguments can be extended to include interactions with frontal cortex in working memory (Frank et al., 2001), as discussed later.

2.0.3 Probabilistic Classification Deficits

Probabilistic classification deficits have been studied using the "weather prediction" task (Knowlton et al., 1994). Participants study sets of cards with multiple cues and have to predict whether the cues presented in a given trial are associated with "rain" or "sunshine". The cue-outcome relationships are probabilistic and not easily determined. Healthy participants implicitly integrate information over multiple trials, progressively improving despite not being able to explicitly state the basis of their choices (Gluck, Shohamy, & Myers, 2002). The BG seems to be recruited for this ability, as it is activated during the learning stages of the weather prediction task (Poldrack, Prabakharan, Seger, & Gabrieli, 1999), and is more generally engaged in tasks that emphasize nondeclarative memory (Poldrack, Clark, PareBlagoev, Shohamy, Moyano, Myers, & Gluck, 2001). The damaged BG in PD likely causes slowed learning observed in patients, just as it has been implicated as a source for habit learning deficits in the motor domain (e.g., Soliveri, Brown, Jahanshahi, Caraceni, & Marsden, 1997; Thomas-Ollivier, Reymann, Le Moal, Schueck, Lieury, & Allain, 1999). But how is the weather prediction task related to habit learning, and exactly what about DA in the BG supports the learning of these so-called habits?

Insight comes from the observation that PD patients are selectively impaired in cognitive procedural learning tasks that involve trial-by-trial error feedback. In purely observational implicit learning tasks (e.g., artificial grammar and prototype learning), patient performance is normative (Reber & Squire, 1999). Among two versions of conditional-associative SR learning, PD patients were only impaired in the one that relied on trial-and-error (Vriezen & Moscovitch, 1990). In implicit categorization tasks, successful integration of information depends on both error feedback (Ashby, Queller, & Berretty, 1999; Ashby, Maddox, & Bohil, 2002) and BG integrity (Ashby, Alfonso-Reese, Turken, & Waldron, 1998).

Taken together, these observations support the notion that feedback mediated learning occurs in the BG and is therefore disrupted in PD. Feedback may modulate DA release in the BG that, in addition to having a performance effect on response execution, is critical for cognitive reinforcement learning.

2.0.4 Phasic Bursting of DA Mediates Trial-and-Error Learning

A healthy range of phasic DA bursts during feedback may lead to the unconscious acquisition of stimulus-reward-response associations. Data reviewed below suggests that positive and negative feedback have opposing effects on DA release, which in turn modulates synaptic plasticity and therefore supports learning.

A multitude of data in primates show that DA releasing cells fire in phasic bursts in response to unexpected reward (Schultz, 1998; Schultz, Dayan, & Montague, 1997). Equally relevant but sometimes ignored, dopaminergic firing dips below baseline when a reward is expected but not received (Hollerman & Schultz, 1998; Schultz, Apicella, & Ljungberg, 1993). In humans, phasic bursts and dips of DA have been inferred to occur during positive and negative feedback, respectively (Holroyd & Coles, 2002).

Several lines of evidence support the notion that these changes in extracellular levels of DA during feedback are critical for learning. First, DA modifies synaptic plasticity in animal experimental conditions. Dopamine D1 receptor stimulation leads to long term potentiation (LTP), whereas D2 stimulation restricts LTP (Nishi, Snyder, & Greengard, 1997). Accordingly, LTP is blocked by D1 antagonists and enhanced by D2 antagonists (for a review, see Centonze, Picconi, Gubellini, Bernardi, & Calabresi, 2001). Second, these effects are behaviorally relevant: administration of D1 antagonists disrupted learning in an appetitive conditioning task, whereas D2 antagonists promoted learning (Eyny & Horvitz, 2003). Third, because dopamine modulates cellular excitability (Nicola, Surmeier, & Malenka, 2000), associative or "Hebbian" learning may be enhanced in the presence of dopamine, since this type of learning depends on the levels of activity of the cells in question (Hebb, 1949; Schultz, 2002). Thus, the efficacy of recently active synapses may be reinforced by a burst of DA acting as a "teaching signal", leading to the learning of rewarding behaviors (Wickens, 1997). This account predicts that a delayed DA burst following the behavior should degrade learning by enhancing the strengths of inappropriate synapses. In human category learning, substantial impairments are indeed observed if feedback is delayed by just 2.5 seconds after each response (Maddox, Ashby, & Bohil, 2003).

In summary, phasic bursts and dips of DA occur differentially during positive and negative feedback, result in modification of synaptic plasticity, and therefore may be critical for the learning of trial-and-error tasks. A plausible explanation for implicit category learning deficits in PD is that damage to dopaminergic neurons in the BG reduces both the tonic and phasic levels of extracellular DA, diminishing the effectiveness of the habit-learning system. Before moving on to a more explicit biologically based version of this theory, the next section discusses the effects of dopaminergic medication on cognition. By artificially increasing levels of DA, medication alleviates some cognitive deficits but actually gives rise to others. This is taken to indicate that the dynamic range of the DA signal may be more critical than its raw level.

2.0.5 Deficits Induced by Dopaminergic Medication

The most common treatments for PD are DA agonists and levodopa (L-Dopa), a DA precursor (Maruyama, Naoi, & Narabayashi, 1996). Many cognitive studies in PD do not take into account the level of medication administered to the patient, somewhat confounding the interpretation of experimental results. That is, if a null effect is found, it could be attributed to the successful replenishment of DA by L-Dopa therapy. Conversely, if an effect is found, it is difficult to know if this effect stems from a lack of DA in PD, or is somehow related to the medication. For instance, medication results in elevated levels of tonic DA in undamaged areas. This may prevent phasic dips from being effective and degrade performance when they are functionally important (e.g., during negative feedback).

A series of studies compared cognitive function in medicated versus non-medicated patients, finding that L-Dopa therapy had positive or deleterious effects on cognitive function, depending on the nature of the task (Gotham et al., 1988; Swainson et al., 2000; Cools et al., 2001). The general conclusion was that dopaminergic medication ameliorates task-switching deficits in PD, but that it impairs performance in **probabilistic reversal** (i.e., learning to reverse stimulus-reward probabilities after prepotent responses are ingrained). Deficits induced by medication are selective to the reversal stage, in which participants must use negative feedback to override

prepotent responses.

The interpretation given by these authors stems from the fact that dopaminergic damage in early stage PD is restricted to the dorsal striatum, leaving the ventral striatum with normal levels of DA (Kish et al., 1988; Agid, Ruberg, Hirsch, Raisman-Vozari, Vyas, Faucheux, Michel, Kastner, Blanchard, Damier, Villares, & Zhang, 1993). This explains why DA medication alleviates deficits in task-switching, which relies on dorsal striatal interactions with dorsolateral prefrontal cortex. However, the amount of medication necessary to replenish the dorsal striatum might "overdose" the ventral striatum with DA, and is therefore detrimental to tasks that recruit it. Reversal learning depends on the ventral striatum and ventral prefrontal cortex in monkeys (Dias, Robbins, & Roberts, 1996; Stern & Passingham, 1995) and recruits these same areas in healthy humans (Cools, Clark, Owen, & Robbins, 2002). The overdose hypothesis is further supported by the finding that medicated, but not non-medicated, patients exhibited impulsive betting strategies in a gambling task known to recruit the ventral striatum (Cools, Barker, Sahakian, & Robbins, 2003).

If the overdose account is accurate, a key question is why should high levels of DA in the ventral striatum produce deficits in reversal learning? Like categorization tasks, reversal learning relies on trial-by-trial feedback. During positive feedback, phasic bursts of DA may still be released. A notable difference is that higher levels of tonic DA might functionally eliminate the effectiveness of phasic dips in DA during negative feedback. A DA agonist would continue to bind to receptors, as it is not modulated by feedback/reward as is endogenous dopamine. This by-product of dopaminergic medication may eliminate an important aspect of the natural biological control system — namely the ability to quickly unlearn previously rewarding behaviors. In non-medicated patients and healthy individuals, phasic dips in DA release may ensue after negative feedback in the reversal stage, allowing the participant to unlearn the prepotent association. The overdose of DA in the ventral striatum of medicated patients would hinder this ability.

So far it has been hypothesized that cognitive deficits in PD arise from a restricted range of DA signals in the BG during error feedback, that does not get completely fixed with medication. This is somewhat vague in that it does not clarify what about the BG supports implicit learning, and how DA modulates processes in the BG. Why should phasic bursts and dips in DA support the learning and unlearning of responses, respectively? To be more clear, I will relate the present discussion to the framework of BG/DA interactions described in Chapter 1, section 1.1. To do so, I first summarize this framework. I then describe a neural network model which instantiates these biological properties and provides a mechanistic account of probabilistic classification and reversal deficits in PD. Besides being a useful tool for understanding complex system interactions in implicit learning, the model can be extended to include those involved in modulating prefrontal function in higher level cognition.

2.0.6 Basal Ganglia/Dopamine Framework

Almost all of the basic mechanisms of the framework put forth in section 1.1 have been postulated in various forms by other researchers. Nevertheless, it represents an integration of these mechanisms into a coherent, mechanistically explicit system. At the most general level, the basal ganglia in this framework modulates the selection of actions being considered in frontal cortex (Mink, 1996; Chevalier & Deniau, 1990; Frank et al., 2001; Frank, in press). More specifically, two main projection pathways from the striatum go through different basal ganglia output structures on the way to thalamus and up to cortex. Activity in the direct pathway sends a "Go" signal to facilitate the execution of a response considered in cortex, whereas activity in the indirect pathway sends a "NoGo" signal to suppress competing responses. Transient changes in dopamine levels that occur during positive and negative feedback have opposite effects on D1 and D2 receptors, which are relatively segregated in the direct and indirect pathways, respectively (Gerfen, 1992; Aubert et al., 2000; Hernandez-Lopez, Bargas, Surmeier, Reyes, & Galarraga, 1997; Hernandez-Lopez, Tkatch, Perez-Garci, Galarraga, Bargas, Hamm, & Surmeier, 2000). Thus the net effect of dopamine bursts during positive reinforcement are to activate the Go pathway and deactivate the NoGo pathway, driving learning so that reinforced responses are subsequently facilitated. Conversely, decreases in dopamine during negative reinforcement have the opposite

effect, driving NoGo learning so that incorrect responses are subsequently suppressed/avoided (Frank, in press).

An important consequence of DA performance effects on Go/NoGo activity levels is that they drive activity-dependent learning to synaptic input. A well established principle should hold across both Go and NoGo cells: more active cells undergo LTP whereas less active cells undergo LTD (e.g., Bear & Malenka, 1994). Once we account for differential effects of DA on excitability in the two BG pathways, this principle makes straightforward predictions on their effects on plasticity. If DA bursts during reinforcement are adaptive, they should have the complementary effects of increasing Go learning while decreasing NoGo learning so that reinforced responses are more likely to be facilitated in the future. Because DA enhances activity in the direct pathway, bursts may indeed induce LTP in Go cells. Further, the inhibitory effects of DA in the indirect pathway may induce LTD in NoGo cells so that they learn to become **less** active. This hypothesis is supported by demonstrations that DA induces LTP via D1 receptors and LTD via D2 receptors (Kerr & Wickens, 2001; Calabresi, Saiardi, Pisani, Baik, Centonze, Mercuri, Bernardi, & Borrelli, 1997).

The same principle can be applied to predict the effect of DA dips, which, if they are adaptive, should enhance NoGo learning so that non-reinforcing responses are actively suppressed in the future. Because DA dips release NoGo cells in the indirect pathway from DA inhibition, the increased NoGo activity should induce LTP in NoGo cells. Although LTP has not been tested during endogenous DA dips, this hypothesis is indirectly supported by examining the effects of D2 receptor blockade, assuming that DA dips decrease D2 stimulation and should therefore have the same qualitative effects on the indirect pathway as D2 blockade. When stimulated by cortical inputs, D2 blockade increases bursting activity and Fos expression of striatal cells in the indirect pathway (Robertson, Vincent, & Fibiger, 1992; Finch, 1999), and results in enhanced corticostriatal LTP (Calabresi et al., 1997).

2.1 Neural Model of BG and DA

The hypothesis is that cognitive deficits in PD can be accounted for by a reduced dynamic range of phasic DA signals which reduces the ability to unconsciously learn Go/NoGo associations. This verbal explanation alone is not sufficient, but may be substantially strengthened by testing its feasibility in a computational model that incorporates all the key elements. Such a model can generate novel predictions because it gets at the underlying source of cognitive dysfunction in PD. If validated, it can also be used as a tool to understand complex involvement of DA in the BG in other neurological disorders.

The above anatomical and biochemical considerations are synthesized in a neural network model (figure 2.1). The model learns to select one of two responses to different input stimuli. Direct and indirect pathways enable the model to learn conditions that are appropriate for gating as well as those for suppressing. Parallel subloops independently modulate each response, allowing selective facilitation of one response with concurrent suppression of the other. Projections from the substantia nigra pars compacta (SNc) to the striatum incorporate modulatory effects of DA. Phasic bursts and dips in SNc firing (and therefore simulated DA release) ensue from correct and incorrect responses, respectively. These phasic changes drive learning by preferentially activating the direct pathway after a correct response and the indirect pathway after an incorrect response. The model is trained on simulated versions of the weather prediction task and probabilistic reversal. Disruption to the DA system as in PD and "overdose" cases produces results that are qualitatively similar to those observed behaviorally.

2.1.1 Mechanics of the Model

The units in the model operate according to a simple **point neuron** function using rate-coded output activations, as implemented in the **Leabra** framework (O'Reilly, 1998; O'Reilly & Munakata, 2000). There are simulated excitatory and inhibitory synaptic input channels. Local inhibition in each of the layers is computed through a simple approximation to the effects of inhibitory interneurons. Synaptic connection weights were trained using a reinforcement learning version of Leabra. The learning algorithm involves two phases, allowing

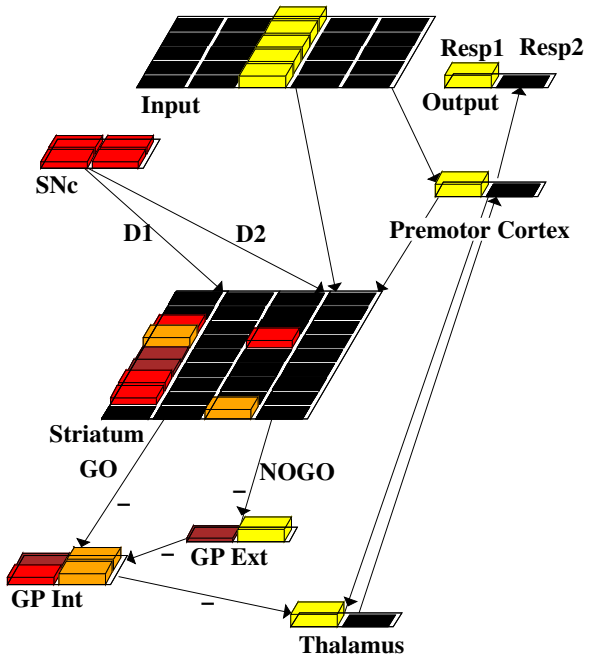

Figure 2.1: Neural network model of direct and indirect pathways of the basal ganglia, with differential modulation of these pathways by DA in the SNc. The Premotor Cortex (PMC) selects a response via direct projections from the Input. BG gating results in bottom-up support from Thalamus, facilitating execution of the response in cortex. In the Striatum, the response has a Go representation (first column) that is stronger than its NoGo representation (third column). This results in inhibition of the left column of GPi and disinhibition of the left Thalamus unit, ultimately facilitating the execution of Response1 in PMC. A tonic level of DA is shown here, during the response selection ("minus") phase. A burst or dip in DA ensues in the feedback ("plus") phase (see figures 2.2 and 2.3), depending on whether the response is correct or incorrect for the particular input stimulus.

simulation of feedback effects, and is more biologically plausible than standard error backpropagation. In the **minus phase**, the network settles into activity states based on input stimuli and its synaptic weights, ultimately "choosing" a response. In the **plus phase**, the network resettles in the same manner, with the only difference being a change in simulated dopamine: an increase for correct responses, and a dip for incorrect responses. Connection weights are then adjusted to learn on the difference between activity states in the minus and plus phases.

2.1.2 Overall Network Division of Labor

The network's job is to select either Response1 or Response2, depending on the task and the sensory input. At the beginning of each trial, incoming stimuli directly activate a response in premotor cortex (PMC). However, these direct connections are not strong enough to elicit a robust response in and of themselves; they also require bottom-up support from the thalamus. The job of the BG is to integrate stimulus input with the dominant response selected by PMC, and based on what it has learned in past experience, either facilitate (Go) or suppress (NoGo) that response.

Within the overall thalamocortical circuit, there are two parallel sub-loops that are isolated from each other, separately modulating the two responses. This allows for the BG to **selectively** gate one response, while continuing to suppress the other(s). This was implemented in our previous model (Frank et al., 2001), and has been suggested by others (Beiser & Houk, 1998). The striatum is divided into two distributed subpopulations. The two columns on the left are "Go" units for the two potential responses, and have simulated D1 receptors. The two columns on the right are "NoGo" units, and have simulated D2 receptors. Thus, the four columns in the striatum represent, from left to right, "Go-Response1"; "Go-Response2"; "NoGo-Response1"; "NoGo-Response2".

The Go columns project only to the corresponding column in the GPi (direct pathway), and the NoGo columns to the GPe (indirect pathway). Both GPe columns inhibit the associated column in GPi, so that striatal Go and NoGo activity have opposing effects on GPi. Finally, each column in GPi tonically inhibits the associated column of the thalamus, which is reciprocally connected to premotor cortex. Thus, if Go activity is stronger than NoGo activity for Response1, the left column of GPi will be inhibited, removing tonic inhibition (i.e. disinhibiting) of the corresponding thalamus unit, and facilitating its execution in premotor cortex.

The above parallel and convergent connectivity is supported by anatomical evidence discussed above. The network architecture simply supports the existence of connections, but how these ultimately influence behavior depends on their relative strengths. The network starts off with random weights and representations in both the BG and cortical layers are learned. Distributed activity within each striatal column enables different Go and NoGo representations to develop for various stimulus configurations during the course of training.

2.1.3 Simulated Effects of DA

To simulate differential effects of DA on D1 and D2 receptors in the two populations of striatal cells, separate excitatory and inhibitory projections were assigned from the SNc to the direct and indirect pathways in the striatum. Thus, the D1 projection only connects to the Go columns of the striatum, whereas the D2 projection connects only to the NoGo columns. Besides being excitatory, the effects of D1 activity involve contrast enhancement. This was accomplished by increasing the striatal units' activation gain (making it more nonlinear), in conjunction with increasing the activation threshold (so that weakly active units do not exceed firing threshold and are suppressed). The effects of D2 activity are inhibitory, suppressing the NoGo cells. Thus, for a high amount of simulated DA, contrast enhancement in the direct pathway supports the enabling of a particular Go response while the indirect pathway is suppressed.

2.1.3.1 DA Modulates Learning

Increases in DA during positive feedback lead to reinforcing the selected response, whereas decreases in DA during negative feedback lead to learning **not** to select that respose. A tonic level of DA is simulated by setting the SNc units to be semi-active (activation value 0.5) at the start of each trial, in the minus phase. In the initial stages of training, the network selects a random response, dictated by random initial weights together with a small amount of random noise in premotor cortex activity. If the response is correct, a phasic increase in SNc firing occurs in the plus phase, with all SNc units set to have an activation value of 1.0 (i.e., high firing rate). This burst of DA causes a more coherent Go representation in the striatum to be associated with the rewarding response that was just selected. For an incorrect response a phasic dip of DA occurs, with all SNc units set to zero activation. In this case, the NoGo cells are released from suppression, enabling the network to learn NoGo to the selected incorrect response.

Note that an explicit supervised training signal is never presented; the model simply learns based on the difference between activity states in the minus and plus phases, which only differ due to phasic changes in DA. Weights from the input layer and premotor cortex are adjusted so that over time, the striatum learns which responses to facilitate and which to suppress in the context of incoming sensory input. In addition, the premotor cortex itself learns to favor a given response for a particular input stimulus, via Hebbian learning from the input layer. Thus, the BG initially learns which response to gate via phasic changes in DA ensuing from random cortical responses, and then this learning transfers to cortex once it starts to select the correct response more often than not. This reflects the idea that the BG is not a stimulus-response module, but rather modulates the gating of responses that are selected in cortex.

2.1.4 Probabilistic Classification Simulations

The weather prediction (WP) task (Knowlton et al., 1994) involves presenting cards made up of four possible cues that have different probabilities of being associated with "rain" or "sun". The predictability of the individual cues is 75.6%, 57.5%, 42.5% and 24.4%. Actual trials involve presenting from one to three cues simultaneously, for a total of 14 cue combinations, making it difficult to become explicitly aware of the probability structure.

In the network, cues are presented in the input, and potential responses are immediately but weakly activated in premotor cortex (figure 2.2). The BG gates one of the two responses if its associated "Go" representation is strong enough, facilitating its execution and suppressing that of the alternative response. If the probabilistically determined feedback to the selected response is positive, a phasic DA burst is applied in the plus phase, resulting in an enhanced Go representation and associated learning. Negative feedback results in a phasic dip of DA in the plus phase, releasing NoGo cells from suppression and allowing the network to learn not to gate the selected response. Over the course of training, networks integrate Go and NoGo signals in the context of different cue combinations to learn when is most appropriate to gate sun and rain responses.

Performance measures involve percentage of **optimal** responses, rather than percentage of responses that were associated with (probabilistically determined) positive feedback provided to the network. Thus, individual responses that had negative outcomes, but were actually the best choice according to the odds, were scored correctly. Similarly, positive outcome responses that were suboptimal were scored incorrectly. These optimal responding measures are consistently used in the behavioral paradigms (e.g., Gluck et al., 2002). Of course, networks were not trained with this error measure but were provided the same probabilistic feedback that would have been given to the human participant.

Further implementational details of the weather prediction task are described in Appendix A.

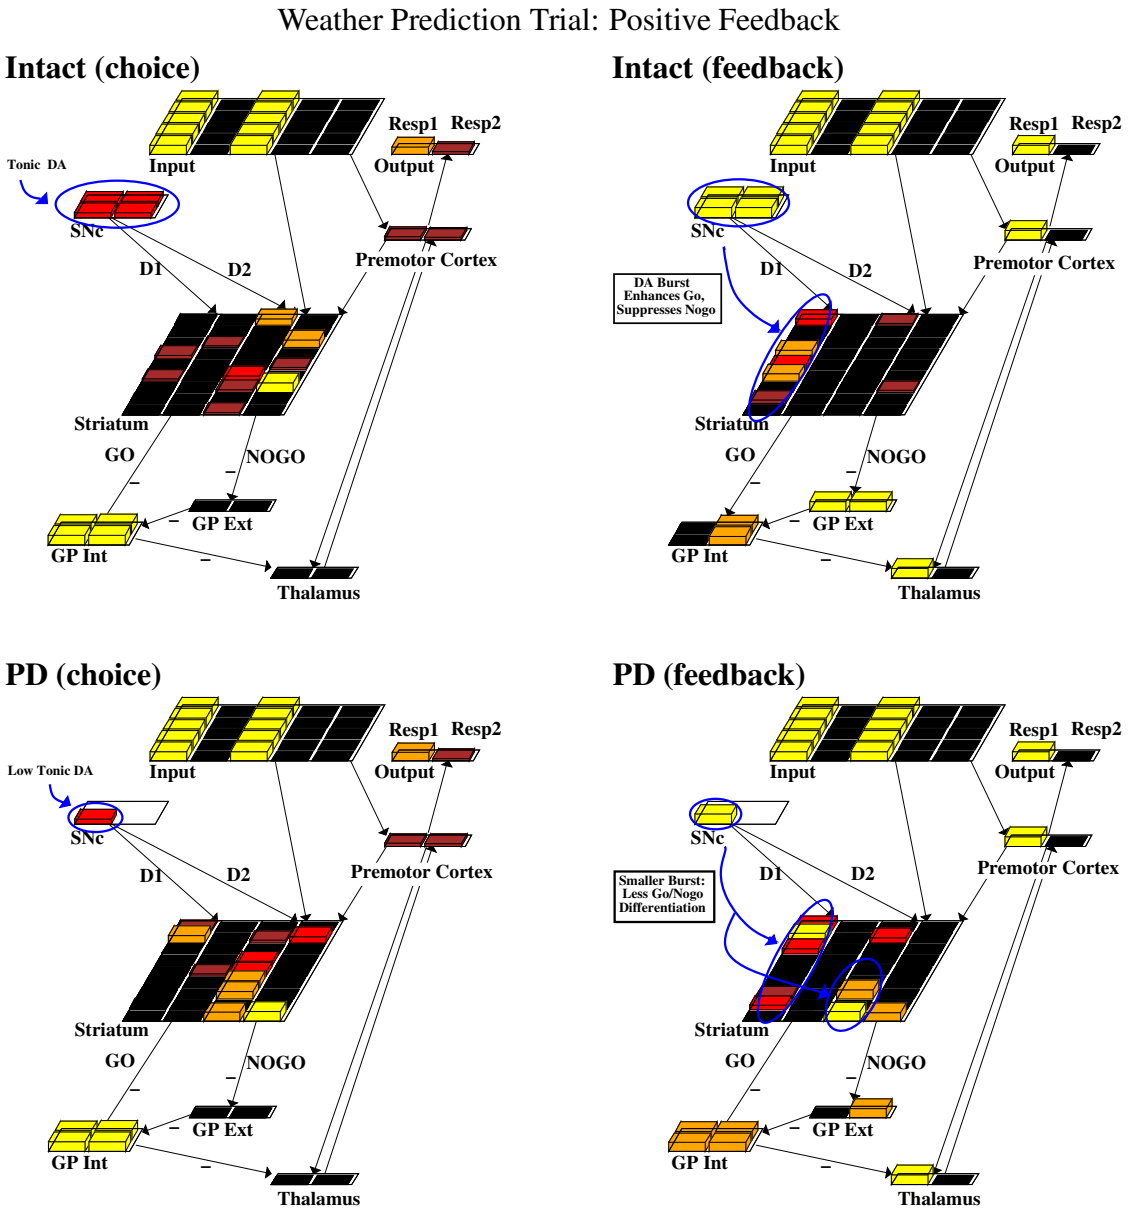

Figure 2.2: A positive feedback trial in the weather prediction task, for both intact and "Parkinson" networks. This trial consists of two cues, represented by the two columns of active units in the Input layer. Intact (choice): Response selection (minus phase) activity in the intact network. Early in training, the BG has not learned to gate either response, as shown by an active GPi and inhibited thalamus. Premotor cortex (PMC) is weakly active due to direct connections from sensory input. The most active (left) unit in PMC, corresponding to "sun", determines the Output response. Intact (feedback): Because the model "guessed" correctly, a phasic burst of DA firing occurs in the SNc, which activates Go units associated with the selected response (via D1 contrast enhancement), while suppressing NoGo units (via D2 inhibition). Weights are adjusted based on differences in network activity between the two phases: The enhanced Go representation during feedback drives learning to gate the "sun" response. PD (choice): Response selection activity in the PD network, for the same trial. Reduced number of intact SNc units results in more active NoGo units. PD (feedback) Correct guessing leads to a phasic burst of DA in the plus phase. However, this phasic burst is not as effective because it applies to only one SNc unit, and therefore only weakly activates more Go units while some NoGo activity persists.

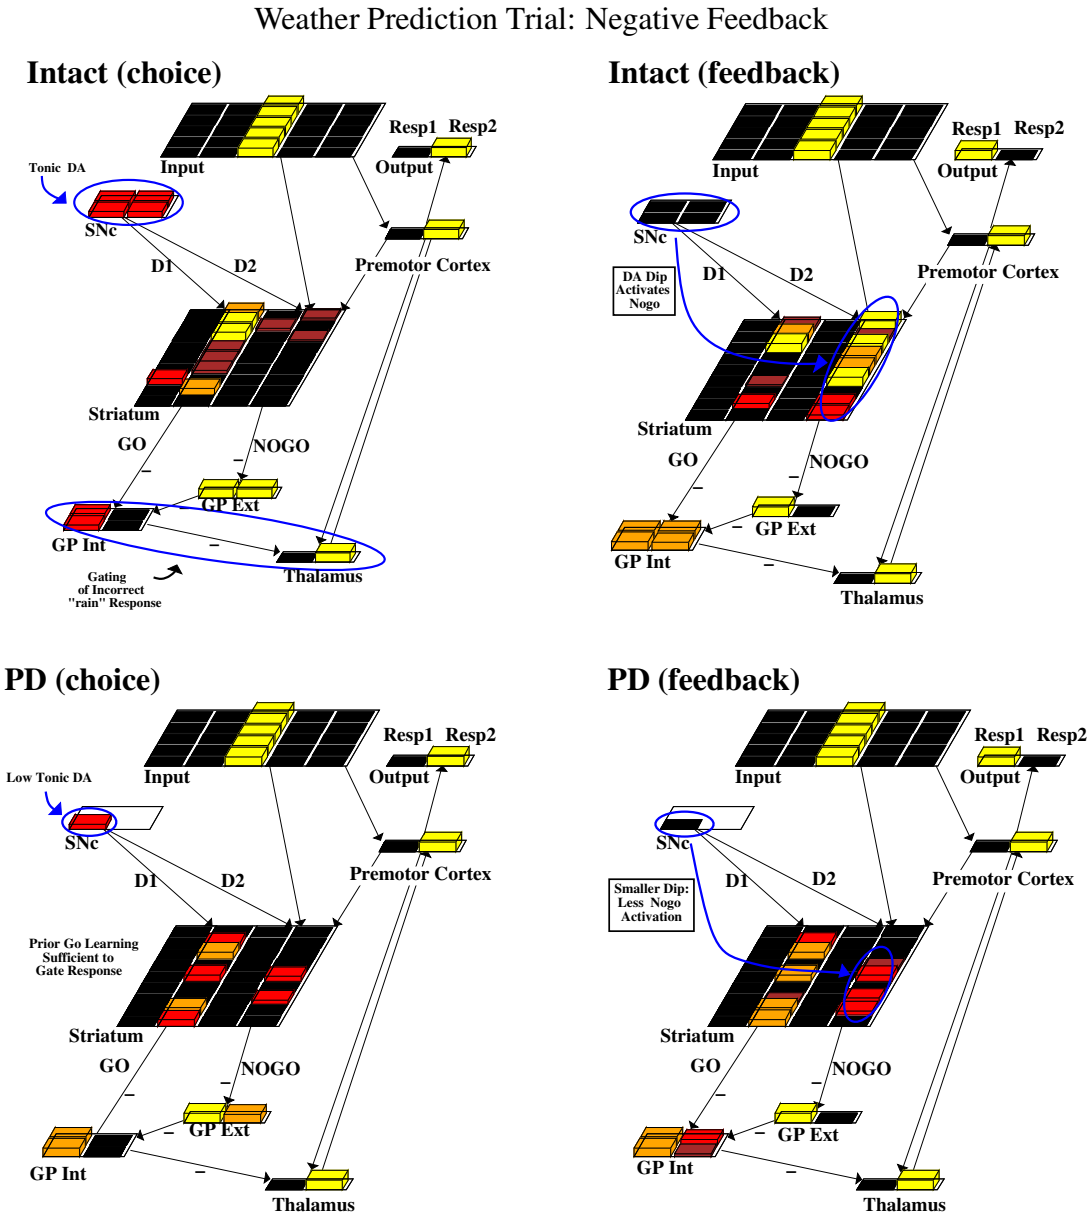

Figure 2.3: A negative feedback trial in the weather prediction task for both intact and "Parkinson" networks. Intact (choice): A single cue is presented. Based on previous learning, the Go units for "rain" are sufficiently active to gate that response, indicated by the inhibition of right GPi units and disinhibition of the right thalamus unit. Intact (feedback): The feedback on this particular trial is negative (due to probabilistic outcomes), shown by a phasic dip of DA firing in the the SNc. The lack of DA removes suppression of NoGo units via D2 receptors, which are then more active than the Go units. The DA dip therefore drives NoGo learning to the incorrect response selected for this cue. Note the Output layer displays the target response for the trial, but this is not used as a training signal: the only signal driving learning is the change in SNc DA. PD (choice): The PD network has also learned to gate the "rain" response for this same trial, based on previous learning. PD (feedback): Feedback is incorrect, and the phasic dip of DA in the SNc leads to activation of some NoGo units. However, the PD network already had low amounts of tonic DA, causing an overall propensity for NoGo learning, so this phasic dip is smaller and therefore not as effective. Reduced dynamic range of DA in the PD network results in less difference in activity levels between the two phases of network settling.

2.1.4.1 Simulated Parkinsonism

Parkinson's disease was simulated by "lesioning" three out of four SNc units so that they were tonically inactive, representing the cell death of approximately 75-80% of dopaminergic neurons in this area (bottom of figures 2.2 and 2.3). This has the effect of reducing tonic DA in the minus phase, as well as phasic DA during feedback in the plus phase. Although the percentage increase/decrease in phasic firing relative to baseline is the same for intact and Parkinson networks, the total amount of DA is reduced by a factor of four, resulting in reduced dynamic range of the DA signal. Dynamic range is critical for learning appropriate Go/NoGo representations from error feedback, as network weights are adjusted based on difference in activity states in the two phases of network settling. NoGo learning is degraded in PD networks because tonic levels of DA are already low, so the phasic dip during negative feedback has less effect. Go learning is degraded because limited amounts of available DA reduce the potency of phasic bursts, activating less of a Go representation during positive feedback.

Less DA in PD also diminishes the contrast enhancement effects of D1 receptor stimulation, further weakening the learning of Go signals. Smaller bursts of DA in PD nets led to less contrast enhancement during positive feedback, by reducing the change in unit activation gain and threshold by a factor of four (see Appendix A). Thus, degraded Go learning is exacerbated because of reduced contrast enhancement that would normally amplify the Go signal during positive feedback.

2.1.4.2 Testing the Contribution of the Indirect Pathway

Since other BG models include the direct, but not necessarily the indirect pathway, the contribution of the latter was evaluated in two different conditions. First, the indirect pathway was disconnected: NoGo units in the striatum no longer projected to the GPe. In these networks, NoGo units were still activated by synaptic input and modulated by DA, but had no effect on BG output. Instead, GPe units tonically inhibited the GPi. This manipulation eliminates the effects of the indirect pathway so that all discrimination learning must be accomplished by comparing Go associations in the direct pathway. Although it is technically possible that this manipulation simply lowers the threshold for gating in the direct pathway by providing more tonic inhibition to GPi (i.e., less overall NoGo), this possibility was accounted for by varying the strength of GPe-GPi inhibitory projections from zero to maximal inhibition. Results reported below are for the best of these cases, which is still substantially worse than the full BG model.

A second test of the indirect pathway was conducted to evaluate the role of response-specific NoGo representations. The hypothesis advocated in this chapter is that each response develops both Go and NoGo representations as a result of positive and negative feedback, and that these representations compete in order to facilitate or suppress the response. However, it is also possible that only Go representations in the striatum are response-specific, and that the indirect pathway represents a more global NoGo signal. In the model, this condition was tested by making GPe units in each column project to both columns of GPi (rather than to just the corresponding column), so that NoGo units in the striatum had the same effect on both responses. Because this may amount to more overall inhibition from GPe to GPi, once again the strength of these inhibitory projections was varied from zero to maximal and the best case results reported.

2.1.5 Probabilistic Reversal Simulation

In the probabilistic reversal (PR) task (Swainson et al., 2000), the participant is presented with two stimuli on a touch-sensitive computer screen and has to choose one of them (by touching it). Feedback provided after each response is probabilistic, with a 80:20% ratio of reinforcement for the 'correct' stimulus. After a number of trials, the probabilities of correct feedback are suddenly reversed, unbeknownst to the participant.

In the model, training involves two stages: acquisition and reversal. In the acquisition stage, the network had to learn which of two stimuli to select. The probabilities associated with correct response were 80% for

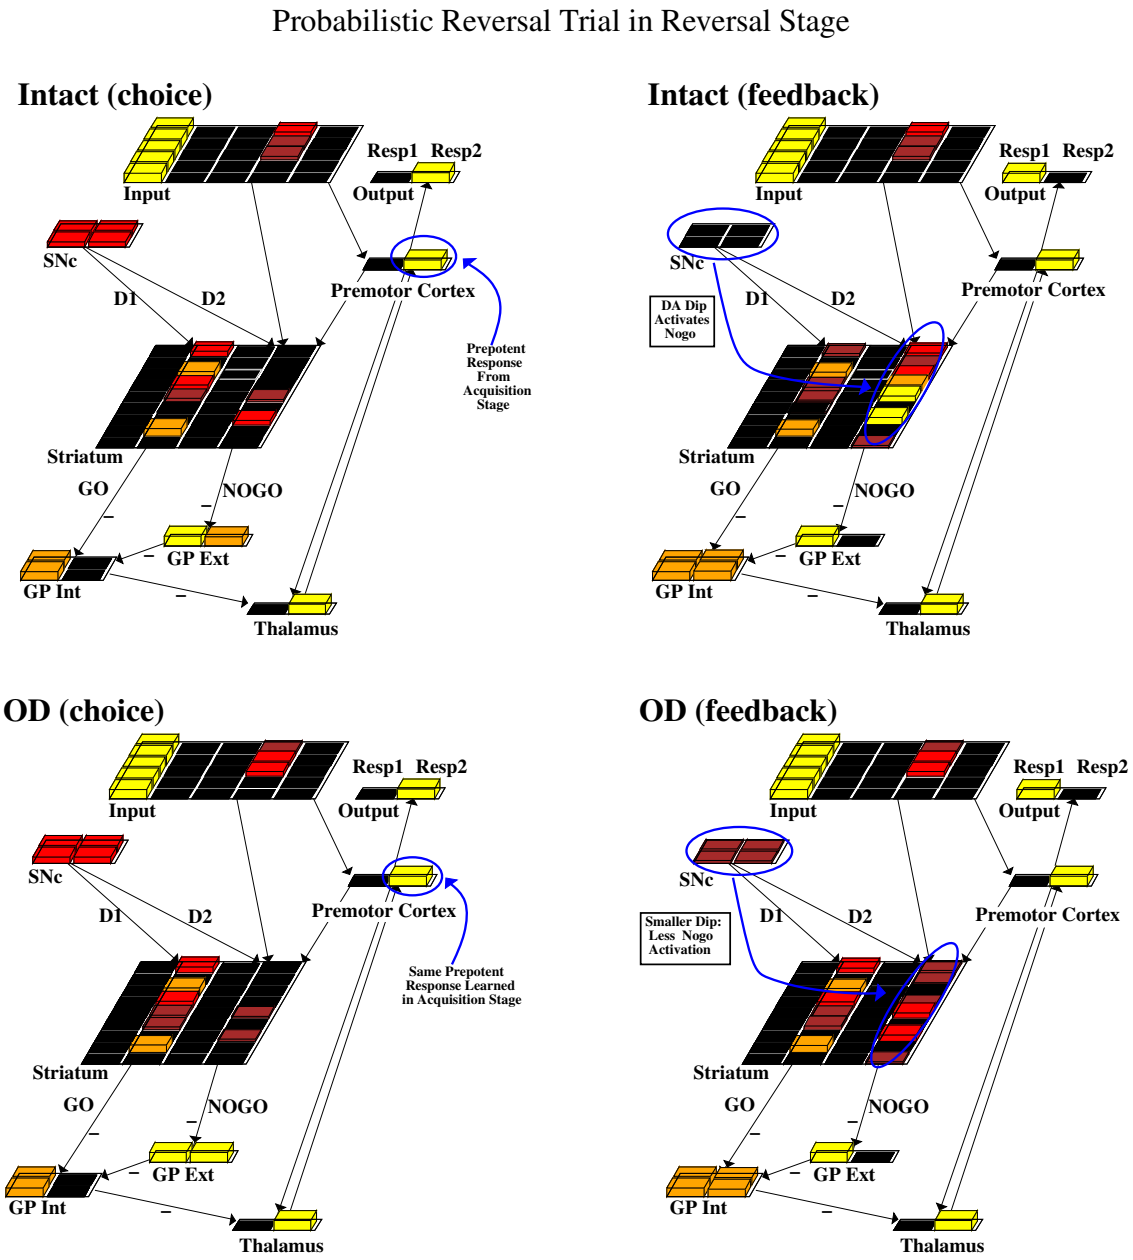

Figure 2.4: A trial in the probabilistic reversal task. Two cues are presented at the input, and the model has to select one of them (see Appendix A for details). The trial shown here is in the reversal stage, during which the model has to learn "NoGo" to the prepotent response before it can switch to selecting the alternative. Reduced dynamic range of DA in the "overdosed" (OD) network causes degraded ability to learn NoGo. Intact (choice): Based on learning in the acquisition stage, the network chooses the stimulus on the right (Resp2). Intact (feedback): In the reversal stage, this choice is incorrect. Phasic dips in SNc release NoGo units from suppression, so that the network can subsequently learn not to perseverate. OD (choice): The same trial is presented to the OD network. Unimpaired Go learning in the acquisition stage results in selection of the same response as the intact network. OD (feedback): A phasic dip is applied to SNc on incorrect trials in the reversal stage, but a residual level of activation due to simulated medication results in weaker activation of NoGo units. The network then takes longer to unlearn the initial response, causing reversal deficits.

selecting stimulus 1 and 20% for selecting stimulus 2. Positive feedback was associated with a DA burst, and negative feedback was associated with a dip. After 50 blocks of trials, these probabilities were reversed, and the feedback effects of DA were necessary to learn NoGo to the prepotent learned response (figure 2.4). Once NoGo representations are strong enough to suppress gating, random cortical activity leads to sometimes choosing the opposite response and DA reinforcement of the corresponding Go representation, enabling reversal.

Further implementational details of the PR task are described in Appendix A.

2.1.5.1 Simulated DA Medication

To model DA overdose in the ventral striatum of medicated patients (Gotham et al., 1988; Swainson et al., 2000; Cools et al., 2001, 2003), all SNc units remained intact. This reflects the fact that the ventral striatum, which is recruited in this task, is relatively spared from dopaminergic damage in moderate PD. The difference between intact and "overdosed" networks was simply an increase in overall level of DA. In the minus phase, the tonic level of DA was increased from an SNc unit activation value of 0.5 to 0.65, reflecting the greater baseline level of DA. Negative feedback in the plus phase resulted in an activation value of 0.25, instead of zero SNc activation. This is still a phasic dip relative to the tonic level, but is meant to simulate the possibility that DA release has less dynamic range in the overdose case (see above for elaboration and justification). Positive feedback resulted in SNc unit activation of 1.0.

The DA overdose manipulations degraded networks' ability to learn NoGo representations during negative feedback, as NoGo units were suppressed by the increased levels of DA. This selectively impairs reversal learning, in which "NoGo" must be learned to a prepotent response.

2.2 Simulation Results

2.2.1 Probabilistic Classification

Despite not having an explicit supervised training signal, simulated phasic DA release during error feedback allowed intact networks to extract the probability structure, scoring 77% optimal responding after 200 trials of training (figure 2.5). "Parkinson" networks were impaired, only scoring 64% optimal responding. Statistical analysis indicated that this difference was highly significant ($F(1,24) = 20.8$, $p = 0.0001$). Two other conditions were run to evaluate the contribution of response-specific NoGo representations in the indirect pathway. Networks with a disconnected indirect pathway were significantly impaired relative to intact networks (65% optimal responding, $F(1,24) = 11.9$, $p = 0.002$). Similarly, networks that had both direct and indirect pathways but only had global NoGo representations (i.e., NoGo units in the striatum affected both responses non-selectively) were also impaired (64% optimal responding, $F(1, 24) = 7.13$, $p = 0.013$). In both these cases, parameters were searched to ensure that impairments were not simply due to an overall threshold for responding, by varying the strength of inhibitory connections from the GPe to the GPi – results reported here are for the best cases over the range from zero to maximal inhibition (which for both cases involved an inhibitory strength of approximately 70% of that in the full model).

2.2.2 Probabilistic Reversal

The results for the PR task were clearcut: "overdosed" networks were selectively impaired at reversing the probabilistic discrimination (figure 2.6). Both intact and overdosed networks were able to acquire the initial 80:20% probabilistic discrimination, with no significant differences between performance in the acquisition phase (97.8 and 98.2 % optimal responding after 200 trials, $F(1, 24) = 0.4$, n.s.). Intact networks consistently learned to reverse this discrimination after a further 200 trials of training, with 78% optimal responding. In contrast, overdosed networks were slower to reverse the initial discrimination, only attaining 64% optimal responding after the same amount of training. These reversal learning differences were significant ($F(1, 24) =$

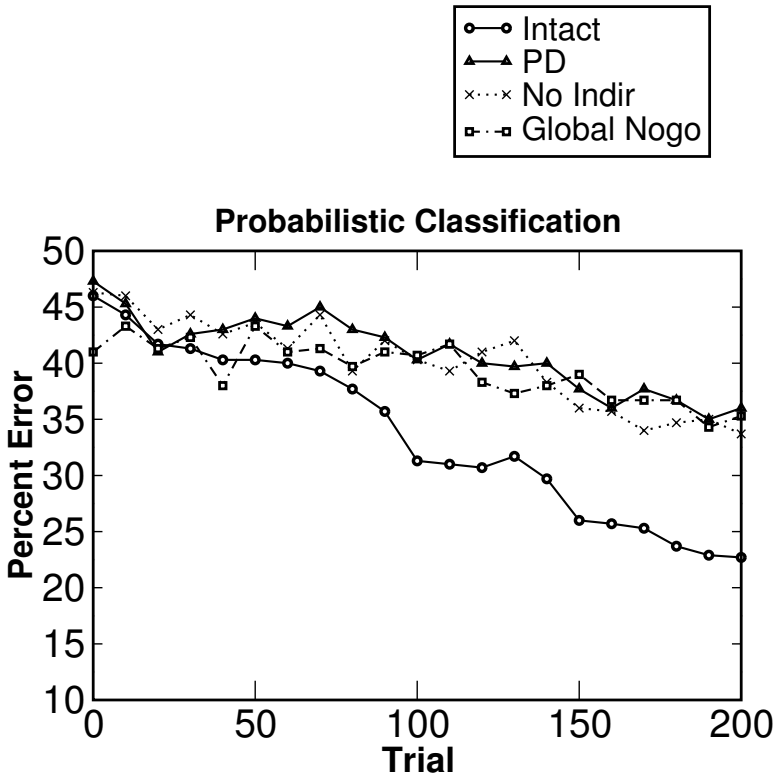

Figure 2.5: Weather prediction task learning curves, averaged over 25 networks for each condition. Intact: Full BG model with direct and indirect pathways modulated by phasic changes in simulated DA during error feedback; PD: simulated Parkinson's disease, modeled by lesioning 75% of dopaminergic units in SNc; No Indir: BG model with the indirect pathway disconnected from the striatum to the GPe; Global NoGo: Full BG model in which NoGo representations globally suppress all responses non-selectively. PD networks are impaired at learning the probabilistic structure, due to impoverished phasic changes in DA in response to feedback. Models without the indirect pathway or with global NoGo representations have reduced discriminability because they can only compare the strength of Go representations to decide which response to facilitate. In contrast, intact models can use response-specific Go and NoGo representations that develop over training in order to more selectively facilitate and suppress responses.

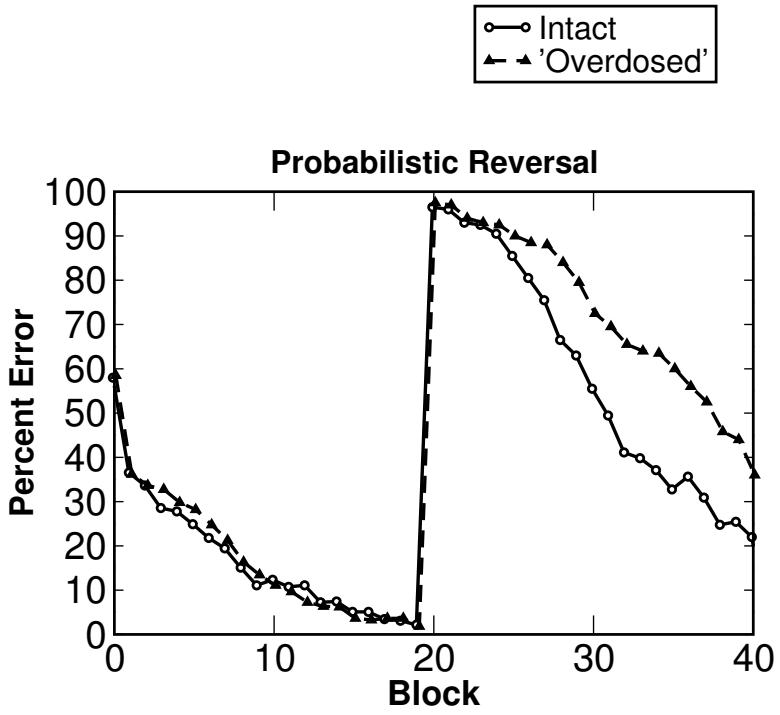

Figure 2.6: Probabilistic reversal results for intact networks and for those with simulated dopamine "overdose", averaged over 25 networks for each condition. Each block consists of 10 trials; reversal of stimulus-outcome probabilities occurred at block 20. Overdosed networks were selectively impaired at learning this reversal, despite performing as well as intact networks in the acquisition phase. A smaller phasic dip in DA during negative feedback resulted in diminished ability to learn "NoGo" to the prepotent response that was learned in the initial acquisition.

4.80, p = 0.038). DA depletion, as is the case for severe PD in the ventral striatum, resulted in non-selective impairment in both stages (not shown).

2.3 Discussion

This work presents a theoretical basis for cognitive procedural learning functions of the basal ganglia. A neural network model incorporating known biological constraints provides a mechanistic account of cognitive deficits observed in PD patients. A key aspect of the model is that phasic changes in dopamine during error feedback are critical for the implicit learning of stimulus-reward-response contingencies, as in probabilistic classification and reversal.

In brief, the model includes competitive dynamics between striatal cells in the direct and indirect pathways of the BG that facilitate or suppress a given response. The cells that detect conditions to facilitate a response provide a "Go" signal, whereas those that suppress responses provide a "NoGo" signal. Habit learning is supported by this circuitry because DA release dynamically modulates the excitability and synaptic plasticity of these pathways so that the most reinforcing responses are subsequently facilitated, while those that are more ambiguous are suppressed.

Simulated Parkinsonism, by reducing the amount of DA in the model and thus its modulatory effects on Go and NoGo representations, produced qualitatively similar results to those observed in PD patients learning the weather prediction task (Knowlton et al., 1994; Knowlton et al., 1996). Less DA led to less contrast enhancement and lower ability to resolve Go/NoGo association differences needed for discriminating between subtly different response reinforcement histories.

Although it could be argued that the simulated indirect pathway is superfluous and that discrimination learning can happen in the direct pathway alone, networks with disrupted indirect pathways were substantially impaired. These results held even when the threshold for facilitating responses in the direct pathway was systematically varied, ensuring that the indirect pathway does not simply set an overall threshold that feasibly could be implemented in the direct pathway alone. Rather, the indirect pathway makes a genuine contribution by developing response-specific NoGo representations that compete with Go representations to enhance discriminability. Without response-specific NoGo representations, the BG is likely to signal "Go" to whichever response happens to be slightly more active in premotor cortex.

2.3.1 Medication-Dependent Deficits

The model also offers some insight as to why patients on medication are impaired at probabilistic reversal. Simulated dopamine overdose produced qualitatively similar results to those observed in medicated patients in probabilistic reversal (Gotham et al., 1988; Swainson et al., 2000; Cools et al., 2001). That is, "overdosed" networks were selectively impaired in the reversal stage, but performed as well or better than control networks in the initial discrimination.

The model provides a mechanistic description of how DA medication may lead to reversal impairments that is generally consistent with the "overdose" hypothesis advocated by authors of the behavioral studies. In intact networks, negative feedback in the reversal stage was associated with phasic dips in DA, which led to activation of striatal NoGo cells by transiently releasing them from the inhibitory influence of DA. The activation of these cells led to suppressing the execution of prepotent responses, allowing networks to learn to reverse responding. In "overdosed" networks, a residual level of DA during negative feedback continued to suppress NoGo cells (via simulated D2 receptors), leading to response perseveration. This account predicts that tonic stimulation of just the D2 receptor should produce similar reversal impairments, which are indeed observed in both healthy humans and non-human primates administered D2 agonists (Mehta, Swainson, Ogilvie, Sahakian, & Robbins, 2001; Smith, Neill, & Costall, 1999).

The above account is also consistent with event-related fMRI studies in humans showing caudate activa-

tion during the reception of negative feedback (Monchi, Petrides, Petre, Worsley, & Dagher, 2001), and ventral striatum activation during the final reversal error in a probabilistic reversal task (Cools et al., 2002). Phasic dips in DA during negative feedback should cause an increase in fMRI signal due to the transient activation of NoGo cells. In the probabilistic reversal task, striatal activation was found specifically during the trial that participants used negative feedback to successfully reverse their behavior on subsequent trials.

Alternative mechanisms are possible to explain reversal learning deficits in patients taking dopaminergic medication. For instance, medication may simply prevent **unlearning** in direct pathway Go cells, rather than suppressing the learning of NoGo cells in the current model. In support of this theory, rats with L-Dopa-induced dyskinesia had a selective impairment in the depotentiation (i.e., reversal of LTP) of corticostriatal synapses (Picconi, Centonze, Hakansson, Bernardi, Greengard, Fisone, Cenci, & Calabresi, 2003), ostensibly due to changes in the D1 receptor pathway. However, it is not clear whether this depotentiation impairment alone can account for reversal deficits: normal depotentiation takes in the order of ten minutes and therefore is not sufficient to induce reversal in a matter of a few trials. Furthermore, the fact that D2 agonists impair reversal implicates a role of the indirect pathway to activate "NoGo" representations and actively avoid situations. Through its push-pull circuitry, the model suggests that the BG is specialized to quickly learn changes in rewarding information.

2.3.2 Relation to Other Models of DA in BG

Other computational models of the BG have focused more on how response selection and reward information may be implemented in biological circuitry (e.g., Brown, Bullock, & Grossberg, 2004a; Gurney, Prescott, & Redgrave, 2001; Beiser & Houk, 1998), but to my knowledge have not attempted to model cognitive implicit learning tasks. Thus it is unclear how prior BG models would account for medicated and non-medicated cognitive impairments in PD. Nevertheless, a comparative analysis of the critical features of the current model with that of others may explicitly demonstrate both consistencies across models as well as novel aspects of the current model that account for behavioral phenomena.

The model builds on earlier work on the interactions between the BG and prefrontal cortex in working memory (Frank et al., 2001), but differs from it in three key aspects. First, the earlier model only included the direct "Go" pathway, as its "NoGo" responses to task-irrelevant stimuli were assumed and hand-wired. The current model includes the competing processes of the indirect pathway, and whether to gate (Go) or suppress (NoGo) a response is learned. Second, the current model includes the SNc/VTA so that the role of dopamine can be implemented, with simulated D1 and D2 receptors in the striatum. Third, the current model does not include the prefrontal cortex or maintenance of information over time, as the simulated tasks do not involve working memory. Instead, the cortical layer in the model is a simpler premotor cortex, representing just two different possible responses (although in principle it could be extended to include several responses).

The model is consistent with other models of DA in the BG (Taylor & Taylor, 1999; Monchi, Taylor, & Dagher, 2000; Brown et al., 2004a), in that dopamine has a performance effect, by differentially modulating excitability in the direct and indirect pathways. A notable difference is that in the current model DA also enhances contrast in the direct pathway by exciting highly active units and suppressing weakly active units, instead of being globally excitatory. This modulatory effect is important for selecting among competing responses, and is motivated by the observed D1 receptor activation excitation of striatal cells in the "up-state", but inhibition of those in the "down-state"(Hernandez-Lopez et al., 1997).

Perhaps a more substantial difference is that while prior models emphasize the tonic effects of DA, the current model also incorporates phasic changes in DA release during positive and negative feedback. Positive feedback results in a phasic burst of DA, transiently biasing the direct pathway and suppressing the indirect pathway. Negative feedback results in a phasic dip in DA, and has the opposite effect. Learning is driven by these transient changes: weight values are modified based on the difference between phases of response selection (hypothesized to involve moderate amounts of DA) and error feedback (hypothesized to involve phasic increases/decreases in DA). Reinforcement learning accounts of DA in the BG have been suggested by others

(e.g., Doya, 2000; Suri & Schultz, 1999), and allow flexible learning of rewarding and non-rewarding behaviors that may change over time.

Another distinguishing feature in the current model is that a particular response is selected by premotor cortex — the BG simply gates this response if it detects the conditions to be appropriate (i.e. predictive of reward). Thus, the BG is not a stimulus-response module, but instead modulates the efficacy of responses being selected in cortex. This is consistent with observations that striatal firing occurs **after** that in premotor cortex and supplementary motor area (e.g., Crutcher & Alexander, 1990; Alexander & Crutcher, 1990b, see also Mink, 1996).[2] Thus, in contrast to the long held assumption that the BG **initiates** motor responses, this model suggests that it facilitates or gates responses that are being considered in premotor cortex. The model further suggests that cortical learning of response selection is mediated by way of DA reward system in the BG, but that once this learning is achieved, the cortex itself selects the response. This is consistent with the observation that Parkinson's patients have specific trouble learning **novel** motor actions, and with the hypothesis that the BG is only necessary for the learning of new categories but not for categorization behavior in experts, which may be mediated directly from perceptual to motor areas (Ashby et al., 1998).

2.3.3 Implications for Frontal Deficits

That PD patients have both implicit learning and frontal deficits — which are not intuitively related — suggests that a better understanding of BG specialization would inform us about how cognition operates as a functional system. The present work only modeled implicit processes in habit learning, which do not have a prefrontal component (Knowlton et al., 1996). However, the same general structure of the model may be extended to include prefrontal cortex, providing insight into the roles of DA and the BG in executive and attentional processes. Indeed, these roles may be very similar to those in implicit learning, with the major difference being the type of representations modulated in the targeted cortical area — motor representations in premotor cortex, and working memory in prefrontal cortex.

Based on the general suggestions of basal ganglia involvement in prefrontal circuits made by Alexander and colleagues (Alexander et al., 1986; Alexander, Crutcher, & DeLong, 1990; Middleton & Strick, 2000b), we developed a computational model that explicitly formulated the role of the BG in working memory (Frank et al., 2001). We suggested that just as the BG facilitates motor command execution in premotor cortex by disinhibiting or "releasing the brakes", it may also facilitate the updating of working memory in prefrontal cortex. If a given stimulus was learned to be task-relevant and therefore suitable for maintenance in PFC, a "Go" signal would be executed by activation of the BG direct pathway, thereby disinhibiting the thalamus and "gating" the updating of PFC.

In the above work (Frank et al., 2001), we briefly discussed the potential role of dopamine, suggesting that it would be important for the learning of task-relevant stimuli via its reward signaling and modulation of synaptic plasticity. In ongoing work (O'Reilly & Frank, in press), we are developing these ideas in a computational model that integrates ventral and dorsal striatum with prefrontal cortex maintenance to demonstrate how complex working memory tasks may be learned (see Chapter 3). Consistent with the present model, DA bursts in the BG preferentially activate cells in the direct pathway via D1 receptors, while suppressing cells in the indirect pathway via D2 receptors. Thus, DA in the BG may have the effect of boosting the updating of working memory by biasing the direct pathway to win the competition for BG output. A phasic dip in DA allows the BG to learn **not** to update task-irrelevant information. The role of DA in the PFC may be quite different, helping to robustly maintain information over time and in the face of interfering stimuli (Durstewitz, Seamans, & Sejnowski, 2000a), depending on optimal levels of DA (Goldman-Rakic, 1996).

With the above model in mind, consider the effect of dopaminergic dysfunction in the BG or PFC. A lack of DA in the BG would lead to too little updating of relevant information into PFC, just as it leads to too

[2] Some, but not all, striatal neurons even fire after the onset of movement. Note that this observation does not conflict with the hypothesized role of the BG to gate or facilitate responses, because firing that occurs after the onset of movement could be associated with either terminating the initiated motor program or suppressing other competing programs.

little execution of motor commands. Conversely, too much DA in the BG would lead to excessive updating of PFC, as observed in L-Dopa induced motor tics and dyskinesia in Parkinson's disease. Finally, a suboptimal level of DA in the PFC would lead to insufficient maintenance of task-relevant information. Any of these DA dysfunctions would lead to "frontal-like" cognitive deficits.

While it is well accepted that the integrity of the PFC is necessary for attentional processes, it is not clear whether attentional deficits seen in PD patients are due to dopaminergic pathology within the PFC itself, or whether DA damage in the BG is sufficient to produce frontal-like deficits due to its interconnections with PFC. In support of the latter possibility, a positive correlation was found between measures of attention and working memory and the level of L-Dopa accumulation in the striatum of PD patients (Remy, Jackson, & Ribeiro, 2000). In monkeys, D2 agents have effects on working memory tasks when applied systemically, but **not** when directly infused into PFC (Granon, Passetti, Thomas, Dalley, Everitt, & Robbins, 2000; Arnsten, Cai, Murphy, & Goldman-Rakic, 1994; Arnsten, Cai, Steere, & Goldman-Rakic, 1995), suggesting that D2 receptors are only indirectly involved in frontal processes.

The current framework holds that D2 effects on working memory are due to modulation of the BG threshold for updating PFC. With high D2 stimulation there is less "NoGo" so the threshold is lowered, and with low D2 stimulation the threshold is raised. Note that a raised threshold means that task-irrelevant stimuli are less likely to get updated. This is consistent with observations that DA depletion to the BG (which should raise the threshold for updating PFC) actually makes monkeys **less** distractible to task-irrelevant stimuli during acquisition of an attentional set (Crofts, Dalley, Collins, Van Denderen, Everitt, Robbins, & Roberts, 2001). However, this higher threshold may also make them more rigid in what to pay attention to, so that they are impaired in task-set switching.

2.3.4 Model Predictions

The main assumption built into the model (supported by data reviewed above) is that positive and negative feedback lead to transient bursts and dips in DA. The model shows that these phasic changes can lead to systems-level effects that modulate the BG threshold for facilitating/suppressing cortical commands. Bursts of DA suppress the NoGo pathway and sharpen representations in the Go pathway. Phasic dips of DA have the opposite effect, releasing the indirect pathway from suppression and allowing the model to learn "NoGo" to the incorrect response. A number of testable predictions can be derived from this model at both neural and behavioral levels.

At the neural level, the model predicts that phasic changes in DA support "Hebbian" learning by modulating neuronal excitability in the indirect pathway via D2 receptors. By transiently suppressing NoGo cells, DA bursts should lead to LTD. Conversely, by transiently exciting NoGo cells, DA dips should lead to LTP. This prediction has not yet been tested directly (during endogenous DA bursts/dips), but is consistent with observations that selective stimulation of D2 receptors leads to LTD, whereas D2 blockade leads to LTP (Calabresi et al., 1997).

The model suggests that a large dynamic range in DA release is necessary for learning subtle differences between positive and negative reinforcement values of responses. Dopamine agonists and antagonists may restrict this range to be at the high and low ends of the DA spectrum, respectively. In a probabilistic reinforcement paradigm, participants administered D2 agonists should easily learn to respond to stimuli having greater than 50% reinforcement probabilities, whereas those taking D2 antagonists (and PD patients) should have an easier time learning to **avoid** stimuli with lower reinforcement probabilities. This is because D2 agonists bias the direct pathway to be more active (by suppressing the indirect pathway), enhancing the learning and execution of Go responses. Parkinson's disease or D2 antagonists should bias the indirect pathway, enabling the learning of NoGo responses.

In addition to modulating the threshold for learning and executing responses, DA should play a similar role in modulating the threshold for updating working memory, discussed in the previous section. D2 agonists should lower this threshold, increasing the amount of updating, whereas D2 antagonists should reduce the

amount of updating. Whether these drugs improve or worsen working memory performance should depend on both the baseline level of updating in the individual (see Kimberg, D'Esposito, & Farah, 1997), and the amount of conflict/interference in the particular task. Specifically, if a working memory task involves distracting information, a lower threshold for updating may result in increased distractibility and impulsiveness because the participant may have trouble ignoring task-irrelevant stimuli. Conversely, if the task simply involves recalling a previously stored memory in the absence of distracting information, D2 agonists should improve performance because they should cause more updating and subsequent maintenance of working memory.

2.3.5 Model Limitations and Future Directions

The model does not differentiate between different parts of the striatum. In fact, the same model is used to simulate probabilistic classification and reversal tasks, which are thought to depend on the dorsal and ventral striatum, respectively. It is at present unclear why these two tasks, which both involve learning response selection via trial-and-error feedback, should involve separate striato-cortical circuits. However, one possibility is that the differences lie in the content of cortical targets: dorsal striatum modulates motor information in premotor cortex, whereas ventral striatum targets reward information in orbitofrontal cortex (OFC) (Alexander et al., 1986; Gottfried, ODoherty, & Dolan, 2003). In reversal learning, a stimulus that has a prepotent reward value suddenly becomes non-rewarding, and OFC representations may be especially important to support top-down activation of "NoGo" representations in the ventral striatum. In probabilistic classification, response selection processes for discriminating among multiple cues may more heavily tax the dorsal striatum. The functional contributions of these two circuits working in tandem will be more explicitly explored in future work.

The present model highlights the importance of dynamic DA modulation in the BG. However, it does not address the brain mechanisms which cause phasic bursts and dips in DA during positive and negative feedback. Instead, this was assumed, and phasic changes in DA were simply set, depending on probabilistic feedback. This implementation does not capture the fact that as learning progresses and rewards become expected, phasic bursts of DA no longer occur during reward but are instead transferred to an earlier stimulus that predicts reward (Ljungberg, Apicella, & Schultz, 1992), as implemented in **temporal differences** (TD) reinforcement learning algorithms (Sutton, 1988). The simple implementation in this model is sufficient for two reasons: a) the tasks are probabilistic, so that positive feedback is never fully predicted (and may therefore always result in a DA burst), and b) even if phasic changes in DA are reduced as outcomes are more expected, this would simply lead to stable performance once the probabilistic structure is learned. A TD-like mechanism is critical in more complex working memory tasks, in which the model has to update and maintain information in one trial to obtain positive reinforcement a few time steps later.

The above discussion makes it clear that the role of DA in the BG in modulating prefrontal function is more complex and needs to be further investigated. While it was briefly considered how DA might modulate the threshold for updating information in PFC, it has not yet been tested whether this would allow for selective updating of task-relevant information, but not that of irrelevant distracting information. By merging aspects of our initial BG-PFC model of working memory (Frank et al., 2001) with that of the current model, we are currently exploring these issues (O'Reilly & Frank, in press), as described in the next chapter.

2.4 Conclusion

When systems level interactions of multiple brain regions are involved, computational investigations provide a valuable complement to experimental brain research. The current model of DA in the BG provides a working hypothesis that can be tested experimentally and behaviorally, as will be evident in chapters 4 and 5.

Chapter 3

Extending the model to working memory and executive function

[1] It is almost universally accepted that the prefrontal cortex plays a critical role in working memory, even though there is little agreement about exactly what working memory is, or how **else** the prefrontal cortex contributes to cognition. Furthermore, it has long been known that the basal ganglia interact closely with the frontal cortex (e.g., Alexander et al., 1986), and that damage to the basal ganglia can produce many of the same cognitive impairments as damage to the frontal cortex (e.g., Brown & Marsden, 1990; Brown, Schneider, & Lidsky, 1997; Middleton & Strick, 2000b). This close relationship raises many questions regarding the cognitive role of the basal ganglia, and how it can be differentiated from that of the frontal cortex itself. Are the basal ganglia and frontal cortex just two undifferentiated pieces of a larger system? Do the basal ganglia and frontal cortex perform essentially the same function, but operate on different domains of information/processing? Are the basal ganglia an evolutionary predecessor to the frontal cortex, with the frontal cortex performing a more sophisticated version of the same function?

We attempted to answer these kinds of questions by presenting a mechanistic theory and implemented computational model of the prefrontal cortex and basal ganglia contributions to working memory (Frank et al., 2001). Specifically, we argued that working memory requires **rapid updating** and **robust maintenance** as achieved by a **selective gating mechanism** (O'Reilly, Braver, & Cohen, 1999; Braver & Cohen, 2000; Cohen, Braver, & O'Reilly, 1996; O'Reilly & Munakata, 2000). Furthermore, although the frontal cortex and basal ganglia are mutually interdependent in our model, we can nevertheless provide a precise division of labor between these systems. On this basis, we can make a number of specific predictions regarding differential effects of frontal vs. basal ganglia damage on a variety of cognitive tasks.

I begin with a brief overview of working memory, highlighting what we believe are the critical functional demands of working memory that the biological substrates of the frontal cortex and basal ganglia must subserve. I show that these functional demands can be met by a selective gating mechanism, which can trigger the updating of some elements in working memory while others are robustly maintained. Building on existing, biologically-based ideas about the basal ganglia role in working memory (e.g., Beiser & Houk, 1998; Dominey, 1995), I show that the basal ganglia are well suited for providing this selective gating mechanism. I then present a neural network model that instantiates our ideas, and performs a working memory task that requires a selective gating mechanism.

3.1 Working Memory

Working memory can be defined as an active system for temporarily storing and manipulating information needed for the execution of complex cognitive tasks (Baddeley, 1986). For example, this kind of memory is clearly important for performing mental arithmetic (e.g., multiplying 42 x 17) — one must maintain subsets of the problem (e.g., 7 x 2) and store partial products (e.g., 14) while maintaining the original problem as well (e.g., 42 and 17) (e.g., Tsung & Cottrell, 1993). It is also useful in problem solving (maintaining and updating

[1] A version of this chapter is published in **Cognitive, Affective and Behavioral Neuroscience** (Frank et al., 2001).

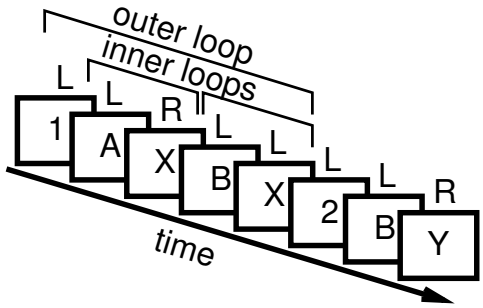

Figure 3.1: The 1-2 CPT-AX task. Stimuli are presented one at a time in a sequence (CPT = continuous performance task), and the subject must respond by pressing the right key (R) to the target sequence, otherwise a left key is pressed. If the subject last saw a 1, then the target sequence is an A followed by an X. If a 2 was last seen, then the target is a B followed by a Y. Distractor stimuli (e.g, 3, C, Z) may be presented at any point in a sequence and are to be ignored. Shown is an example sequence of stimuli and the correct responses, emphasizing the inner- and outer-loop nature of the memory demands (maintaining the task stimuli (1 or 2) is an outer-loop, while maintaining the prior stimulus of a sequence is an inner-loop).

goals and subgoals, imagined consequences of actions, etc), language comprehension (keeping track of many levels of discourse, using prior interpretations to correctly interpret subsequent passages, etc), and many other cognitive activities (see Miyake & Shah, 1999 for a recent survey).

From a neural perspective, one can identify working memory with the maintenance and updating of information encoded in the active firing of neurons (**activation-based memory**) (e.g., Fuster, 1997; Goldman-Rakic, 1987). It has long been known that the prefrontal cortex exhibits this kind of sustained active firing over delays (e.g., Fuster & Alexander, 1971; Kubota & Niki, 1971; Miyashita & Chang, 1988; Funahashi, Bruce, & Goldman-Rakic, 1989; Miller, Erickson, & Desimone, 1996). Such findings support the idea that the prefrontal cortex is important for active maintenance of information in working memory.

3.1.1 Working Memory Functional Demands

The A-X version of the Continuous Performance Task (CPT-AX) is a standard working memory task that has been extensively studied in humans (Cohen, Perlstein, Braver, Nystrom, Noll, Jonides, & Smith, 1997; Braver & Cohen, 2000). The subject is presented with sequential letter stimuli (A, X, B, Y), and is asked to detect the specific sequence of an A followed by an X by pushing the right button. All other combinations ($A-Y$, $B-X$, $B-Y$) should be responded to with a left button push. This task requires a relatively simple form of working memory, where the prior stimulus must be maintained over a delay until the next stimulus appears, so that one can discriminate the target from non-target sequences. We have devised an extension of this task that places somewhat more demands on the working memory system. In this extension, which we call the 1-2-AX task (figure 3.1), the target sequence varies depending on prior **task demand** stimuli (a 1 or 2). Specifically, if the subject last saw a 1, then the target sequence is $A-X$. However, if the subject last saw a 2, then the target sequence is $B-Y$[2] . Thus, the task demand stimuli define an **outer loop** of active maintenance (maintenance of task demands) within which there can be a number of **inner loops** of active maintenance for the A-X level sequences.

The full 1-2-AX task places three critical functional demands on the working memory system:

Rapid updating: As each stimulus comes in, it must be rapidly encoded in working memory (e.g., one-trial updating, which is not easily achieved in weight-based memory).

[2] Other variations in target sequences for the two sub-tasks are possible, and are being explored empirically.

Robust maintenance: The task demand stimuli (1 or 2) in the outer loop must be maintained in the face of interference from ongoing processing of inner loop stimuli and irrelevant distractors.

Selective updating: Only some elements of working memory should be updated at any given time, while others are maintained. For example, in the inner loop, A's and X's (etc) should be updated while the task demand stimulus (1 or 2) is maintained.

One can obtain some important theoretical leverage by noting that the first two of these functional demands are directly in conflict with each other, when viewed in terms of standard neural processing mechanisms (Cohen et al., 1996; Braver & Cohen, 2000; O'Reilly et al., 1999; O'Reilly & Munakata, 2000). Specifically, rapid updating can be achieved by making the connections between stimulus input and working memory representations strong, but this directly impairs robust maintenance, as such strong connections would allow stimuli to interfere with ongoing maintenance. This conflict can be resolved by using an active **gating** mechanism (Cohen et al., 1996; Hochreiter & Schmidhuber, 1997).

3.1.2 Gating

An active gating mechanism dynamically regulates the influence of incoming stimuli on the working memory system. When the gate is open, stimulus information is allowed to flow strongly into the working memory system, thereby achieving rapid updating. When the gate is closed, stimulus information does not strongly influence working memory, thereby allowing robust maintenance in the face of ongoing processing. The computational power of such a gating mechanism has been demonstrated in the LSTM model of Hochreiter and Schmidhuber (1997), which is based on error backpropagation mechanisms and has not been related to brain function, and in more biologically-based models by Braver and Cohen (2000) and O'Reilly and Munakata (2000).

These existing biologically-based models provide the point of departure for the present model. These models were based on the idea that the neuromodulator **dopamine** can perform the gating function, by transiently strengthening the efficacy of other cortical inputs to the frontal cortex. Thus, when dopamine release is phasically elevated, as has been shown in a number of neural recordings (e.g., Schultz et al., 1993), working memory can be updated. Furthermore, these models incorporate the intriguing idea that the same factors that drive dopamine spikes for learning (e.g., Montague, Dayan, & Sejnowski, 1996) should also be appropriate for driving working memory updating. Specifically, working memory should be updated whenever a stimulus triggers an enhanced prediction of future reward. However, an important limitation of these models comes from the fact that dopamine release is relatively global — large areas of prefrontal cortex would therefore receive the same gating signal. In short, a dopamine-based gating mechanism does not support the **selective** updating functional demand listed above, where some working memory representations are updated as others are being robustly maintained. Therefore, the present model explores the possibility that the basal ganglia can provide this selective gating mechanism, as described in the previous chapters for the gating functions of the basal ganglia in lower level cognitive tasks.

In short, one can think of the overall influence of the basal ganglia on the frontal cortex as "releasing the brakes" for motor actions and other functions. Put another way, the basal ganglia are important for **facilitating** motor movements, but not for determining the detailed properties of these movements (e.g., Hikosaka, 1989; Chevalier & Deniau, 1990; Passingham, 1993; Bullock & Grossberg, 1988). Clearly, this disinhibitory gating in the motor domain could easily be extended to gating in the working memory domain. Indeed, this suggestion was made by Chevalier and Deniau (1990) in generalizing their ideas from the motor domain to the cognitive one. Subsequently, several theories and computational models have included variations of this idea (Alexander et al., 1990; Goldman-Rakic & Friedman, 1991; Dominey & Arbib, 1992; Houk & Wise, 1995; Dominey, 1995; Gelfand, Gullapalli, Johnson, Raye, & Henderson, 1997; Beiser & Houk, 1998). Thus, we find a striking convergence between the functionally-motivated gating ideas we presented earlier and similar ideas developed more from a bottom-up consideration of the biological properties of the basal ganglia/frontal cortex system.

Specifically, in the context of the working memory functions of the frontal cortex, our model is based on the idea that the basal ganglia are important for **initiating the storage of new memories**. In other words, the disinhibition of the thalamocortical loops by the basal ganglia results in the opening of the gate into working memory, resulting in rapid updating. In the absence of striatal firing, this gate remains closed, and the frontal cortex maintains existing information. Critically, the basal ganglia can provide a **selective** gating mechanism because of the many parallel loops. Although the original neuroanatomical studies suggested that there are around 5 such loops (Alexander et al., 1986), it is likely that the anatomy can support many more subloops within these larger-scale loops (e.g., Beiser & Houk, 1998), meaning that relatively fine-grained selective control of working memory is possible.

To summarize, at least at this general level, it appears that the basal ganglia can provide exactly the kind of selective gating mechanism that our functional analysis of working memory requires. Our detailed hypotheses regarding the selective gating mechanisms of this system are specified in the following sections.

3.1.3 Details of Active Maintenance and the Gating Mechanism

I begin with a discussion of the mechanisms of active maintenance in the frontal cortex, which then constrain the operation of the gating mechanism provided by the basal ganglia.

Perhaps the most obvious means of achieving the kinds of actively maintained neural firing observed in prefrontal cortex neurons using basic neural mechanisms is to have **recurrent excitation** among frontal neurons resulting in **attractor states** that persist over time (e.g., Dehaene & Changeux, 1989; Zipser, Kehoe, Littlewort, & Fuster, 1993; Seung, 1998; Moody, Wise, di Pellegrino, & Zipser, 1998; Braver & Cohen, 2000; O'Reilly & Munakata, 2000). With this kind of mechanism, active maintenance is achieved because active neurons will provide further activation to themselves, perpetuating an activity state. Most theories/models of the basal ganglia role in working memory employ a variation of this type of maintenance, where the recurrent connections are between frontal neurons and the thalamus and back (Hikosaka, 1989; Alexander et al., 1990; Goldman-Rakic & Friedman, 1991; Dominey & Arbib, 1992; Houk & Wise, 1995; Dominey, 1995; Gelfand et al., 1997; Beiser & Houk, 1998; Taylor & Taylor, 2000). This form of recurrence is particularly convenient for enabling the basal ganglia to regulate the working memory circuits, as thalamic disinhibition would directly facilitate the flow of excitation through the thalamocortical loops.

However, it is unclear if there are sufficient numbers of thalamic neurons relative to frontal neurons to support the full space of maintainable frontal representations. When a given thalamic neuron sends activation to the frontal cortex to support maintenance, its connectivity would have to uniquely support one particular representation or part thereof, otherwise the specificity of the maintained information will be lost. Therefore, the number of thalamic neurons would have to be on the same order as that of the frontal neurons, unless frontal representations are massively redundant. Recurrent connectivity within the frontal cortex itself avoids this problem. Furthermore, we are not aware of any definitive evidence suggesting that these loops are indeed critical for active maintenance (e.g., showing that frontal active maintenance is eliminated with selective thalamic lesions, which is presumably a feasible experiment). Another issue with thalamocortically-mediated recurrent loops is that they would generally require persistent disinhibition in the thalamus during the entire maintenance period (though see Beiser & Houk, 1998 for a way of avoiding this constraint). For these reasons, we are inclined to think in terms of intracortical recurrent connectivity for supporting frontal maintenance.

Although it is intuitively appealing, the recurrence-based mechanism has some important limitations stemming from the fact that information maintenance is entirely dependent on the instantaneous activation state of the network. For example, it does not allow for the frontal cortex to exhibit a transient, stimulus-driven activation state and then return to maintaining some previously encoded information — the set of neurons that are most active at any given point in time will receive the strongest excitatory recurrent feedback, and will therefore be what is maintained. If a transient stimulus activates frontal neurons above the level of previously maintained information, then this stimulus transient will displace the prior information as what is maintained.

This survival-of-the-most-active characteristic is often violated in recordings of prefrontal cortex neu-

rons. For example, Miller et al. (1996) observed that frontal neurons will tend to be activated transiently when irrelevant stimuli are presented while monkeys are maintaining other task relevant stimuli. During these stimulus transients, the neural firing for the maintained stimulus can be weaker than that for the irrelevant stimulus. After the irrelevant stimuli disappear, the frontal activation reverts to maintaining the task-relevant stimuli. We interpret this data as strongly suggesting that frontal neurons have some kind of **intrinsic** maintenance capabilities.[3] This means that individual frontal neurons have some kind of intracellular "switch" that, when activated, provides these neurons with extra excitatory input that enhances their capacity to maintain signals in the absence of external input. Thus, this extra excitation enables maintaining neurons to recover their activation state after a stimulus transient — after the actual stimulus ceases to support its frontal representation, the neurons with intrinsic excitation will dominate.

There are a number of possible mechanisms that could support a switchable intrinsic maintenance capacity for frontal neurons (e.g., Lewis & O'Donnell, 2000; Fellous, Wang, & Lisman, 1998; Wang, 1999; Dilmore, Gutkin, & Ermentrout, 1999; Gorelova & Yang, 2000; Durstewitz, Seamans, & Sejnowski, 2000b). For example, Lewis and O'Donnell (2000) report clear evidence that, at least in an anesthetized preparation, prefrontal neurons exhibit bistability — they have **up** and **down** states. In the **up** state, neurons have a higher resting potential and can easily fire spikes. In the **down** state, the resting potential is more negative, and it is more difficult to fire spikes. A number of different possible mechanisms are discussed by Lewis and O'Donnell (2000) that can produce these effects, including selective activation of excitatory ion channels in the **up** state (e.g., Ca^{2+} or Na^{+}), or selective activation of inhibitory K^{+} ion channels in the **down** state.

Other mechanisms that involve intracellular switching, but depend more on synaptic input, have also been proposed. These mechanisms take advantage of the properties of the NMDA receptor, which is activated both by synaptic input and by postsynaptic neuron depolarization, and produces excitation through Ca^{2+} ions (Fellous et al., 1998; Wang, 1999; Durstewitz, Kelc, & Gunturkun, 1999; Durstewitz et al., 2000a). In the model by Wang and colleagues, a switchable bistability emerges as a result of interactions between NMDA channels and the balance of excitatory and inhibitory inputs. In the model by Durstewitz and colleagues, dopamine modulates NMDA channels and inhibition to stabilize a set of active neurons, and prevent interference from other neurons (via the inhibition). Consistent with these models, we think that recurrent excitation plays an important maintenance role in addition to a switchable intrinsic maintenance capacity. As we discuss below, recurrent excitation can provide a "default" maintenance function, and it is also important for magnifying and sustaining the effects of the intrinsic maintenance currents.

There are many complexities and unresolved issues with these maintenance mechanisms. For example, although dopamine clearly plays an important role in some of these mechanisms, it is not clear if tonic levels present in awake animals would be sufficient to enable these mechanisms, or whether phasic bursts of dopamine would be required. This can have implications for the gating mechanism, as we discuss later. Despite the tentative nature of the empirical evidence, there are enough computational advantages of a switchable intrinsic maintenance capacity (as combined with a more conventional form of recurrent excitation), to compel us to use such a mechanism in our model. Furthermore, we think the neurophysiological finding that working memory neurons recover their memory-based firing even after representing transient stimuli (as reviewed above) makes a compelling empirical case for the presence of such mechanisms.

There are two primary computational advantages to a switchable intrinsic maintenance capacity. The first is that it imparts a significant degree of robustness on active maintenance, as has been documented in several models (e.g., Fellous et al., 1998; Durstewitz et al., 2000a). This robustness stems from the fact that intrinsic signals are not dependent on network dynamics, whereas spurious strong activations can hijack recurrent maintenance mechanisms. Second, these intrinsic maintenance mechanisms, by allowing frontal cortex to represent both transient stimuli and maintained stimuli, avoid an important catch-22 problem that arises in bootstrapping learning over delays (O'Reilly & Munakata, 2000) (figure 3.2). Briefly, learning that it is useful to maintain

[3] Although it is still possible that other frontal areas were really maintaining the signal during the intervening stimulus activations, this explanation becomes less appealing as this phenomenon is consistently observed across many different frontal areas.

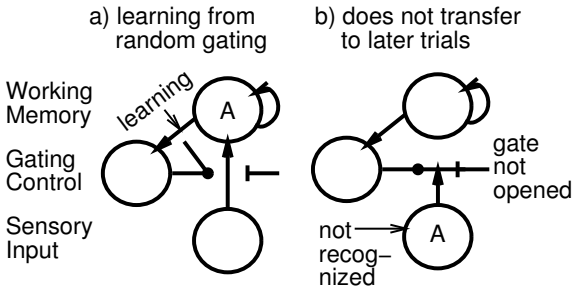

Figure 3.2: Illustration of the catch-22 problem that occurs when the gating mechanism learns based on maintained working memory representations, and those representations can only become activated after the gating mechanism fires for a given stimulus. a) Learning about a stimulus A presented earlier and maintained in frontal cortex, which is based on initially random exploratory gating signals, will be between the maintained representations and the gating controller. b) When this stimulus is later presented, it will not activate the working memory representations until the gate is opened, but the gate has only learned about this stimulus from these same working memory representations, which are not activated.

a stimulus can only occur after that stimulus has been maintained in frontal representations, meaning that the gating mechanism must learn what to maintain based on frontal representations. However, if these frontal representations only reflect stimuli that have already been gated in for maintenance, then the gating mechanism will not be able to detect this stimulus as something to gate in until it is already gated into frontal cortex! However, if the frontal representations always reflect current stimuli as well as maintained information, then this problem does not occur.

Dynamic gating in the context of an intracellular maintenance switch mechanism amounts to the activation and deactivation of this switch. Neurons that participate in the maintenance should have the switch turned on, and those that do not should have the switch turned off. This contrasts with other gating models developed in the context of recurrent activation-based maintenance, which required gating to modulate the strength of input weights into the frontal cortex (e.g., Braver & Cohen, 2000; O'Reilly & Munakata, 2000), or the strength of the thalamocortical recurrent loops (e.g., Dominey, 1995; Gelfand et al., 1997; Beiser & Houk, 1998). Therefore, we propose that the disinhibition of the thalamocortical loops by the basal ganglia results in the modulation of the intracellular switch. Specifically, we suggested that the activation of the layer 4 frontal neurons that receive the excitatory projection from the thalamus (or the equivalent cell types in layer 3 in motor areas of frontal cortex — we just use the layer 4 notation for convenience) is responsible for modulating intracellular ion channels on the neurons in other layers (which could be in either layers 2-3 or 5-6) that are ultimately responsible for maintaining the working memory representations.

Finally, more conventional recurrent excitation-based maintenance is important in our model for establishing a "default" propensity of the frontal cortex to maintain information. Thus, if nothing else has been specifically gated on in a region of frontal cortex (i.e., if no other neurons have a specific competitive advantage due to intracellular maintenance currents), then the recurrent connectivity will tend to maintain representations over time anyway. However, any new stimulus information will easily displace this kind of maintained information, and it cannot compete with information that has been specifically gated on. This default maintenance capacity is important for "speculative" trial-and-error maintenance of information — the only way for a learning mechanism to discover if it is important to maintain something is if it actually does maintain it, and then it turns out to be important. Therefore, having a default bias to maintain is useful. However, this default maintenance bias is overridden by the active gating mechanism, allowing learning to have full control over what is ultimately maintained.

To summarize, in our model, active maintenance operates according to the following set of principles:

- Stimuli generally activate their corresponding frontal representations when they are presented.
- Robust maintenance occurs only for those stimuli that trigger the intracellular maintenance switch (as a result of the conjunction of external excitation from other cortical areas and layer 4 activation resulting from basal ganglia-mediated disinhibition of the thalamocortical loops).
- When other stimuli are being maintained, those representations that did not have the intracellular switch activated will decay quickly following stimulus offset.
- However, if nothing else is being maintained, recurrent excitation is sufficient to maintain a stimulus until other stimuli are presented. This "default" maintenance is important for learning by trial-and-error what is relevant to maintain.

3.1.4 Additional Anatomical and Computational Constraints

In this section, I discuss the implications of a few important anatomical properties of the basal ganglia/frontal cortex system. First, I consider consequences of the relative sizes of different regions in the basal-ganglia frontal cortex pathway. Next, I examine evidence that can inform the number of different separately gatable frontal areas. Finally, I discuss the level of convergence and divergence of the loops, and then raise several anatomical properties that are not current incorporated into our model.

A strong constraint on understanding basal ganglia function comes from the fact that the GPi and SNr have a relatively small number of neurons — there are approximately 111 million neurons in the human striatum (Fox & Rafols, 1976), whereas there are only 160,000 in the GPi (Lange, Thorner, & Hopf, 1976) and a similar number in the SNr. This means that whatever information is encoded by striatal neurons must be vastly compressed or eliminated on its way up to the frontal cortex. This constraint coincides nicely with the gating hypothesis — the basal ganglia do not need to convey detailed **content** information to the frontal cortex — instead they simply need to tell different regions of the frontal cortex **when** to update. As noted in the context of motor control, damage to the basal ganglia appears to affect **initiation**, but not the details of **execution** of motor movements — presumably not that many neurons are needed to encode this gating or initiation information.

Given this dramatic bottleneck in the GPi/SNr, one might wonder why there are so many striatal neurons in the first place. We think this is also sensible under the gating proposal: in order for only task-relevant stimuli to get updated (or an action initiated) via striatal firing, these neurons need to only fire for a very specific **conjunction** of environmental stimuli and internal context representations (as conveyed through descending projections from frontal cortex). This context-specificity of striatal firing has been established empirically (e.g., Schultz, Apicella, Romo, & Scarnati, 1995), and is an important part of many extant theories/models (e.g., Wickens, 1993; Houk & Wise, 1995; Wickens, Kotter, & Alexander, 1995; Berns & Sejnowski, 1998; Jackson & Houghton, 1995; Beiser & Houk, 1998; Amos, 2000). Thus, many striatal neurons are required to encode all of the different specific conjunctions that can be relevant. Without such conjunctive specificity, there would be a risk that striatal neurons would fire for inappropriate subsets of stimuli. For example, the 1 and 2 stimuli should be maintained separately from the other stimuli in the 1-2-AX task, but this is not likely to be true of other tasks. Therefore, striatal neurons should encode the conjunction of the stimulus (1 or 2) together with some representation of the 1-2-AX task context from the frontal cortex. If the striatum instead employed a smaller number of neurons that just respond to stimuli without regard to task context (or other similar kinds of conjunctions), confusions between the many different implications of a given stimulus would result. Note that by focusing on conjunctivity in the striatum we do not mean to imply that there is no conjunctivity in the frontal representations as well (e.g. Watanabe, 1992; Rao, Rainer, & Miller, 1997) — frontal conjunctive representations can be useful for maintaining appropriately contextualized information.

Another constraint to consider concerns the number of different subregions of the frontal cortex for which the basal ganglia can plausibly provide separate gating control. Although it is impossible to determine any precise estimates of this figure, even the very crude estimates provided here are informative in suggesting that gating occurs at a relatively fine-grained level. Fine-grained gating is important for mitigating conflicts where two representations require separate gating control, and yet fall within one gating region. An upper limit estimate is provided by the number of neurons in the GPi/SNr, which is roughly 320,000 in the human as noted previously. This suggests that the gating signal operates on a **region** of frontal neurons, instead of individually controlling specific neurons (and, assuming the thalamic areas projecting to frontal cortex are similarly sized, argues against the notion that the thalamocortical loops themselves can maintain detailed patterns of activity).

An interesting possible candidate for the regions of frontal cortex that are independently controlled by the basal ganglia are distinctive anatomical structures consisting of interconnected groups of neurons, called **stripes** (Levitt, Lewis, Yoshioka, & Lund, 1993; Pucak, Levitt, Lund, & Lewis, 1996). Each stripe appears to be isolated from the immediately adjacent tissue, but interconnected with other more distal stripes, forming a cluster of interconnected stripes. Furthermore, it appears that connectivity between the prefrontal cortex and thalamus exhibits a similar, though not identical, kind of discontinuous stripe-like structuring (Erickson & Lewis, 2000, 2004). Therefore, it would be plausible that each stripe or cluster of stripes constitutes a separately controlled group of neurons — each stripe can be separately updated by the basal ganglia system. Given that each stripe is roughly .2-.4 mm by 2-4 mm in size (i.e., .4-1.6 mm^2 in area), one can make a rough computation that the human frontal cortex (having roughly 1/4 of the approximate 140,000 mm^2 surface area of the entire cortex; Douglas & Martin, 1990) could have over 20,000 such stripes (assuming that the stripes found in monkeys also exist in humans, with similar properties). If the thalamic connectivity were with stripe clusters and not individual stripes, this figure would be reduced by a factor of around 5. In either case, given the size of the GPi

and SNr, there would be some degree of redundancy in the per-stripe gating signal at the GPi/SNr level. Also note that the 20,000 (or 4,000 for stripe clusters) figure is for the entire frontal cortex, so the proportion located in the prefrontal cortex (and thus involved in working memory function) would be smaller. Further evidence consistent with the existence of such stripe-like structures comes from the finding of iso-coding microcolumns of neighboring neurons that all encode roughly the same information (e.g., having similar directional coding in a spatial delayed response task) (Rao, Williams, & Goldman-Rakic, 1999).

The precise nature of the inputs and outputs of the loops through the basal ganglia can have implications for the operation of the gating mechanism. From a computational perspective, it would be useful to control each stripe using a range of different input signals from the sensory and frontal cortex (i.e., broad convergence of inputs), to make the gating appropriately context-specific. In addition, it is important to have input from the current state of the stripe that is being controlled, as this would affect whether this stripe should be updated or not. This implies closed loops going through the same frontal region. Data consistent with both of these connectivity patterns has been presented (see Graybiel & Kimura, 1995; Middleton & Strick, 2000a for reviews). Although some have taken mutually exclusive positions on these two patterns of connectivity, and the facts are a matter of considerable debate, both patterns are mutually compatible from the perspective of our model. Furthermore, even if it turns out that the cortical projections to the striatum are relatively focused, context sensitivity in gating can be achieved via context-sensitive frontal input representations. In other words, the context sensitivity of gating could come either from focused context-sensitive inputs to the striatum, or from broad sensory inputs that are integrated by the striatum itself. One particularly intriguing suggestion is that the convergence of inputs from other frontal areas may be arranged in a hierarchical fashion, providing a means for more anterior frontal areas (which may represent higher-level, more abstract task/goal information) to appropriately contextualize more posterior areas (e.g., supplementary and primary motor areas) (Gobbel, 1997).

To summarize, anatomical constraints are consistent with the selective gating hypothesis by suggesting that the basal ganglia interacts with a large number of distinct regions of the frontal cortex. We hypothesize that these distinct stripe structures constitute separately-gated collections of frontal neurons, extending the parallel loops concept of Alexander et al. (1986) to a much finer grained level (see also Beiser & Houk, 1998). Thus it is possible to maintain some information in one set of stripes, while **selectively** updating other stripes.

3.1.5 The Motor Control-Working Memory Continuum

We have emphasized that our view of the basal ganglia interactions with frontal cortex builds on existing ideas regarding these interactions in the context of motor control. Specifically, both the initiation of a motor act and the updating of working memory representations require striatal firing to disinhibit or gate frontal cortex representations. Although we have discussed motor control and working memory as two separable functions, it is probably more useful to think in terms of a continuum between cognitive working memory and motor control functions. For example, one can think of the neurons in premotor or supplementary motor areas as maintaining a motor control plan that guides a sequence of basic motor movements (e.g., Shima & Tanji, 1998; Wise, 1985b). This plan would need to be maintained over the duration of the sequence, and can thus be considered a working-memory representation. Thus, the line between working memory and motor control is fuzzy — indeed, this ambiguity provides useful insight as to why both motor control and working memory are co-localized within the frontal cortex.

3.1.6 Summary: The Division of Labor between Frontal Cortex and Basal Ganglia

Before describing our model in detail, and by way of summary, we return to the fundamental question posed at the outset of this paper — what is the nature of the division of labor between the frontal cortex and the basal ganglia? In light of all the foregoing information, we offer the following concise summary of this division of labor:

- The frontal cortex uses continuously-firing activations to encode information over time in working

memory (or, on a shorter time scale, to execute motor actions).

- The basal ganglia fires only at very select times to trigger the updating of working memory states (or initiate motor actions) in frontal cortex.

Furthermore, we can speculate as to **why** it would make sense for the brain to have developed this division of labor in the first place. Specifically, one can see that the use of continuously-firing activation states to encode information is at odds with the need to only fire at very specific times. Therefore, the brain may have separated these two systems to develop specialized mechanisms supporting each. For example, striatal neurons must have a relatively high effective threshold for firing, and it can be difficult to regulate such a threshold to ensure that firing happens when appropriate, and not when it is not appropriate. The dopaminergic neuromodulation of these neurons, and its control by descending projections from the striatum, may be important specializations in this regard. Finally, we do not mean to claim that all striatal neurons only fire in a punctate manner — others exhibit sustained "delay period" activations (e.g., Alexander, 1987; Schultz et al., 1995). We think these reflect sustained frontal activations, not an intrinsic maintenance capability of striatal neurons themselves. Similarly, punctate firing in cortical neurons, especially in motor output areas, could be a reflection of gating signals from the basal ganglia.

3.2 The 1-2-AX Model

We have implemented the ideas outlined above in a computational model of the 1-2-AX task. This model demonstrates how the basal ganglia can provide a selective gating mechanism, by showing that the outer-loop information of the task demand stimuli (1 or 2) can be robustly maintained while the inner loop information (A, B etc) is rapidly updated. Furthermore, we show that irrelevant distractor stimuli are ignored by the model, even though they transiently activate their frontal representations. In addition, the model demonstrates that the same mechanisms that drive working memory updating also drive the motor responses in the model.

3.2.1 The Mechanics of the Model

The model is shown in Figure 3.3. The units in the model operate according to a simple **point neuron** function using rate-coded output activations, as implemented in the **Leabra** framework (O'Reilly & Munakata, 2000; O'Reilly, 1998). There are simulated excitatory synaptic input channels, and inhibitory input is computed through a simple approximation to the effects of inhibitory interneurons. There is also a constant leak current that represents the effects of K^+ channels that are always open, and the maintenance frontal neurons have a switchable excitatory ion channel that is off by default. See the appendix for the details and equations. The model's representations were predetermined, but the specific weights were trained using the standard Leabra error-driven and associative (Hebbian) learning mechanisms to achieve target activations for every step in the sequence. In a more realistic model, the representations would not be predetermined but rather would develop as a function of the learning mechanisms — this shortcut was simply used as a convenient way of achieving a desired set of representations to test the basic sufficiency of our ideas about the gating mechanism.

For simplicity, every layer in the model has been organized into three different "stripes," where a stripe corresponds to an individually updatable region of frontal cortex, as discussed previously. The right-most stripe in each layer represents the outer-loop task demand information (1 or 2). The middle stripe represents information maintained at the inner-loop, sequence-level (A or B). The left-most stripe represents stimuli that actually trigger an action response (X or Y). To clarify and simplify the motor aspects of the task, we only have a response at the end of an inner-loop sequence (i.e., after an X or Y), instead of responding L for all the preceding stimuli. All these other other responses should be relatively automatic, whereas the response after the X or Y requires taking into account all the information maintained in working memory, so it is really the task-critical motor response.

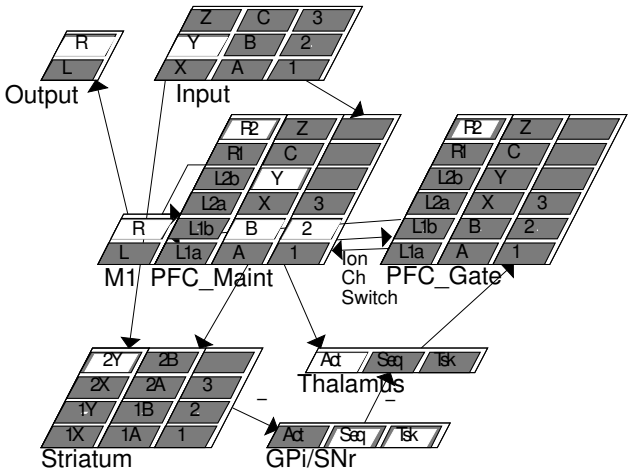

Figure 3.3: Working memory model with basal ganglia mediated selective gating mechanism. The PFC has been subdivided into maintenance (PFC_Maint) and gating (PFC_Gate) layers. Three hierarchically organized "stripes" of the PFC and basal ganglia are represented as the three columns of units within each layer — each stripe is capable of being independently updated. The right-most **task** stripe encodes task-level information (i.e., 1 or 2). The middle **sequence** (seq) encodes sequence-level information within a task (i.e., A or B). The left-most **action** (act) stripe encodes action-level information (i.e., responding to the X or Y stimulus and actually producing the left or right output in PFC). Non task-relevant inputs (e.g., $3, C, Z$) are also presented and the model ignores them, i.e., they are not maintained.

We describe the specific layers of the model in the course of tracing a given trial of input. First, a stimulus is presented (activated) in the Input layer. Every stimulus automatically activates its corresponding frontal representation, located in the PFC_Maint layer of the model. This layer represents cortical layers 2-3 and 5-6 (without further distinguishing these layers, though it is possible there are divisions of labor between them), and is where stimulus information is represented and maintained. The other frontal layer is PFC_Gate, which represents the gating action of cortical layer 4 — we'll return to it in a moment.

If the input stimulus has been recognized as important for task performance, as a result of as-yet unimplemented learning experiences (which are represented in the model through hand-set enhanced weight values), then it will activate a corresponding unit in the Striatum layer. This activation of the high-threshold striatal unit is the critical step in initiating the cascade of events that leads to maintaining stimuli in working memory, via a process of "releasing the brakes" or disinhibiting the thalamic loops through the frontal cortex. Note that these striatal units in the model encode **conjunctions** of maintained information in frontal cortex (1 or 2 in this case) and incoming stimulus information (A, B, X, or Y). Although not computationally essential for this one task, these conjunctions reflect our theorizing that striatal neurons need to encode conjunctions in a high-threshold manner to avoid task-inappropriate stimulus activation. The frontal representations are also necessarily conjunctive in their detection of the combination of stimuli that trigger a response action — the stimulus maintenance representations could also be more conjunctive as well, even though it is not strictly necessary for this one task.

Once a striatal unit fires, it inhibits the globus pallidus unit in its corresponding stripe, which has to this point been tonically active and inhibiting the corresponding thalamus unit. Note the compression of the signal from the striatum to the globus pallidus, as discussed above. The disinhibition of the thalamic unit opens up the recurrent loop that flows from the PFC_Maint units to the thalamus and back up to the PFC_Gate layer. Note that the disinhibited thalamic unit will only get activated if there is also descending activation from PFC_Maint units. Although this is always the case in our model, it wouldn't be true if a basal ganglia stripe got activated (disinhibited) that did not correspond to an area of frontal activation — this property may be important for synchronizing frontal and basal ganglia representations during learning.

The effect of thalamic firing is to provide general activation to an entire stripe of units in the PFC_Gate layer. These frontal units cannot fire without this extra thalamic activation, but they also require excitation from units in the PFC_Maint layer, which are responsible for selecting the specific gate unit to activate. Although this is configured as a simple one-to-one mapping between maintenance and gating frontal units in the model, the real system could perform important kinds of learning here to fine-tune the gating mechanism. Finally, the activation of the gating unit controls the switchable excitatory ion channels in the frontal maintenance units. For those maintenance units within a stripe that receive both input from the current input stimulus and the gating activation, the excitatory ion channels are opened. Maintenance units that only get the gating activation, but not stimulus input, have their ion channels closed. This mechanism provides a means of updating working memory by resetting previously active units that are no longer receiving stimulus input, while providing sustained excitatory support for units that do have stimulus input.

3.2.2 An Example Sequence

Figure 3.4 shows an example sequence of $2-B-C-Y$ as processed by the model. The first stimulus presented is the task context — in this case it is task 2, the $B-Y$ detection task. Because the striatum detects this stimulus as being task relevant (via the 2 striatal unit), it inhibits the task globus pallidus unit, which then disinhibits the corresponding thalamus unit. This disinhibition enables the thalamus to then become excited via descending projections from frontal cortex. The thalamic activation then excites the PFC_Gate unit that also receives activation from the PFC_Maint layer, resulting in the activation of the excitatory ion channel for the 2 frontal unit in the PFC_Maint layer.

Next, the B input activates the $2B$ conjunctive striatal unit, which detects the combination of the 2 task maintained in frontal cortex and the B stimulus input. This results in the firing of the sequence stripe

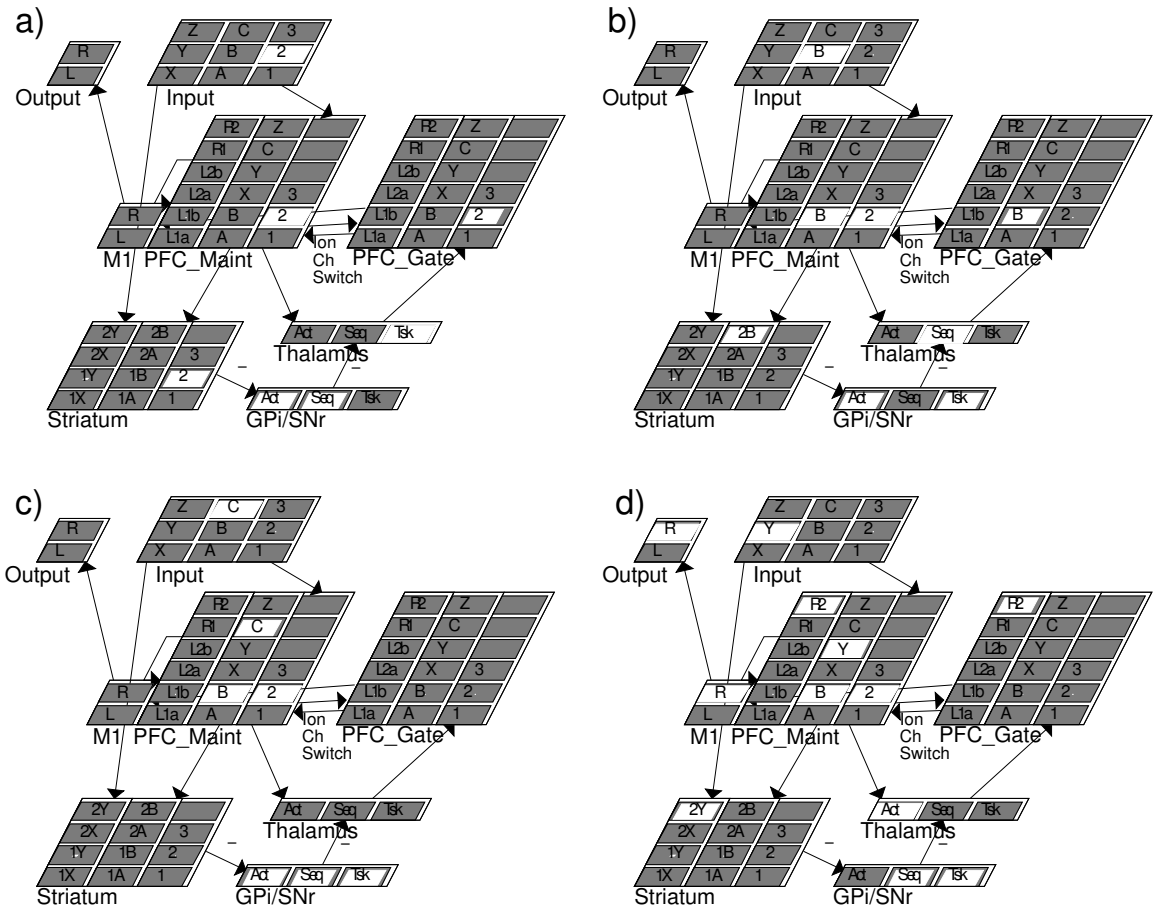

Figure 3.4: An example sequence in the model ($2-B-C-Y$). a) Task context 2 is presented. The striatum detects this stimulus as relevant and disinhibits the task stripe of the thalamus, allowing PFC_Gate to become active, causing the task number to be maintained in PFC_Maint. b) The next stimulus is B, which the striatum detects in conjunction with task context 2 (from the PFC) via the $2B$ unit. The sequence stripe of the thalamus is then disinhibited and B is gated into PFC_Maint, while task context 2 remains active due to persistent ionic currents. This demonstrates **selective** gating. c) A distractor stimulus C is presented, and because the striatum has not built up relevant associations to this stimulus, all units are sub-threshold. The thalamus remains inhibited by the tonically active globus pallidus, and C is not maintained in the PFC. d) Stimulus Y is presented, and the striatum detects the conjunction of it and the task context via the $2Y$ unit. The thalamus action level stripe is disinhibited, which activates conjunctive units in the frontal cortex ($R2$) that detect combinations of maintained and input stimuli ($2-B-Y$). These frontal units then activate the R response in the primary motor area (M1).

and maintenance of the B stimulus encoding in frontal cortex. Note that the 2 has been maintained as the B stimulus was processed and encoded into active memory, due to the fact that these items were represented in different stripes in the frontal cortex. This demonstrates the principle of selective gating, which is central to our model.

The next stimulus is a C distractor stimulus — this is not detected as important for the task by the striatum (i.e., all striatal units remain sub-threshold), and is thus not gated into robust active maintenance (via the intrinsic ion channels). Note that despite this lack of gating, the C representation is still activated in the PFC_Maint frontal cortex layer, as long as the stimulus is present. However, when the next stimulus comes in (the Y in this case), the C activation decays quickly away.

Finally, the Y stimulus is important because it triggers an action. The $2Y$ striatal unit enables firing of the $R2$ unit in the PFC layers — this is a conjunctive unit that detects the conjunction of all the relevant working memory and input stimuli ($2 - B - Y$ in this case) for triggering one kind of R output response (the other R conjunction would be a $1-A-X$). This conjunctive unit then activates the basic R motor response, in a manner consistent with observed frontal recordings (e.g., Hoshi, Shima, & Tanji, 2000). Thus, the same basal-ganglia mediated disinhibitory function supports both working memory updating and motor response initiation in this model.

Although it is not represented in this example, the model will maintain the 2 task signal over many inner-loop sequences (until a different task input is presented), because the inner-loop updating is selective and therefore does not interfere with maintenance of the outer-loop task information.

3.2.3 Summary

To summarize, the model illustrates how frontal cortex can maintain information for "contextualizing" motor responses in a task appropriate fashion, while the basal ganglia trigger the updating of these frontal representations, and the initiation of motor responses.

3.3 A Model of Learning to Update Prefrontal Working Memory Representations

Implicit in our gating model is that the basal ganglia somehow know when it is appropriate to update working memory representations. More recently, we have addressed the "homunculus problem" by merging the BG/DA model described in the previous chapter, which has both Go and NoGo pathways that are modulated by DA during learning, with the simpler working memory model just described (O'Reilly & Frank, in press). The more complex BG-PFC-DA model learns how to update prefrontal cortex (PFC) working memory representations to solve complex cognitive tasks (O'Reilly & Frank, in press). This model (Figure 3.5), called the prefrontal-basal ganglia working memory model (PBWM), also includes an implementation of the neural systems that control the firing of the midbrain DA neurions in the VTA and SNc, which in turn drive the learning of the dorsal striatal neurons that control updating of PFC working memory representations. These DA-based learning mechanisms are based on both bottom-up biological considerations, and top-down computational considerations for achieving successful learning of complex working memory tasks. Our work on this model to date has focused on a few basic tasks, such as the 1-2-AX task (Frank et al., 2001; Kroger et al., in preparation). This task requires multiple levels of working memory maintenance, including an "outer loop" of maintaining the current task goal (i.e., look for an A-X sequence or a B-Y sequence), and an "inner loop" of maintaining the last stimulus that was seen (e.g., an A or B) so that the target sequence can be properly identified. The PBWM model can learn, based on DA bursts generated entirely by the model in response to correct/incorrect reward feedback, to fire the appropriate pattern of BG Go and NoGo signals to selectively update different inner and outer loop representations in PFC. We have also tested it on a range of other challenging tasks.

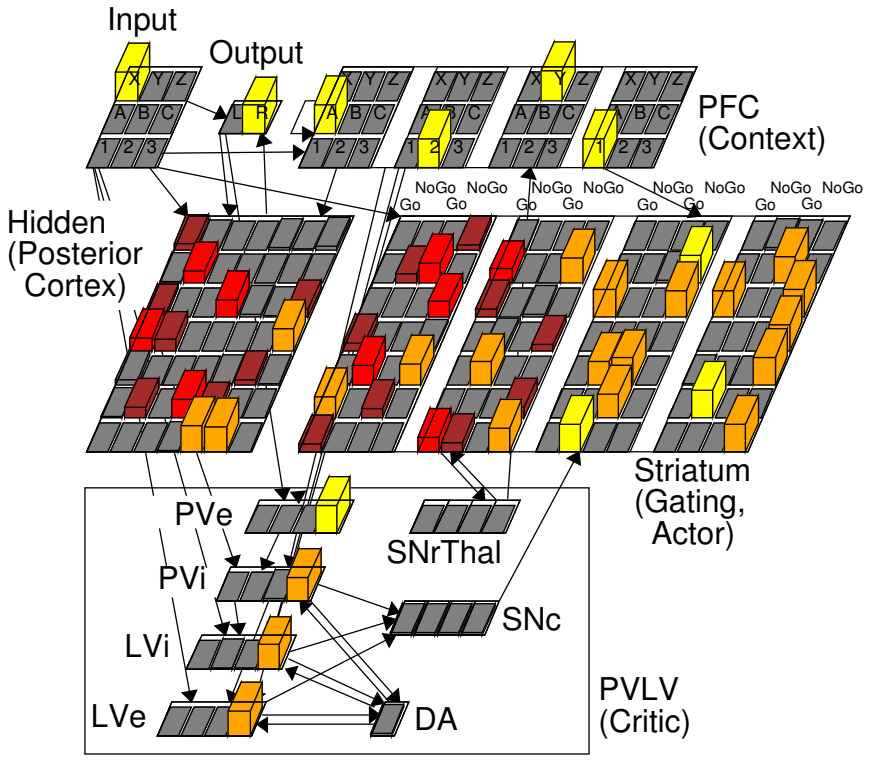

Figure 3.5: Implemented PBWM model as applied to the 1-2-AX task. There are 4 independently-updatable "stripes" through the PFC/BG system in this model, as indicated by the groups of units within the PFC and BG Matrix (and the 4 units in the SNc and SNrThal layers). ExtRew represents external reward (r). NAcLV, NAcPV and ABL together form the "Pavlovian" system that learns to predict r, VTA and SNc compute the DA values from these PVLV layers , and SNc projects this modulation to the Matrix. Go and NoGo units alternate (from bottom-left to upper-right) in the Matrix. The SNrThal layer represents the output of the BG; firing in each unit dictates updating and subsequent maintenance (via intrinsic ion currents) of the corresponding PFC stripe. PFC provides context for Input/Hidden/Output mapping areas, which represent posterior cortex.

3.4 Extending the Framework: New Directions

I have argued that the BG selectively facilitates the most appropriate or potentially rewarding motor action or working memory representation being considered in frontal cortex, while suppressing less appropriate representations from causing interference. Furthermore, I have shown how the modulatory effects of dopamine can explain a number of behavioral findings in Parkinson's disease, and other populations. In the studies described later, I describe experiments that have successfully tested some detailed predictions derived from this framework. But despite these important advances, considerable work remains to be done. In particular, further investigation is needed to more specifically delineate what the BG can and cannot do by itself.

For instance, patients with ventromedial prefrontal damage are known to make suboptimal decisions (i.e., they don't maximize long-term rewards) (Bechara et al., 1998). If the BG can dissociate rewarding commands from less rewarding ones, then why can't these patients, whose BG are intact, make appropriate decisions? One possiblility is that phasic changes in DA that occur during positive and negative reinforcement do not have sufficient information to indicate different overall **values** of rewards, but only encode their relative frequencies (Frank & Claus, in preparation). The ventromedial and/or orbitofrontal cortices may be critical for representing relative magnitudes of reward (e.g., Tremblay & Schultz, 1999) in an active state that is required for top-down decision making. The recruitment of OFC may be especially essential when the expected value of a particular decision conflicts with its relative frequency—that is, when long-term optimal decisions have to compete against other decisions that would lead to more frequent, but smaller overall rewards. Similarly, prefrontal involvement may be instrumental for making choices that lead to large rewards in the future at the expense of achieving small rewards in the present. Preliminary simulations of BG-OFC interactions have successfully captured this ideas, showing specific impairments of OFC lesions in conditions when the expected value of a choice is in opposition to the relative frequency of it resulting in positive outcomes (Frank & Claus, in preparation).

The ventral striatum (VS) is densely interconnected with the OFC and anterior cingulate (Groenewegen et al., 1999). Our framework holds that, similar to the dorsal striatum gating of working memory into dorsolateral PFC, the ventral striatum gates reward information to be maintained in OFC (Frank, under revisions). The VS may also play a critical role in selecting frontal actions (motor responses and updating of OFC reward representations), depending on situational **context**. Indeed, the nucleus accumbens (a principal component of the VS) is often thought of as the interface between limbic desires and motor output behavior, because it integrates context information in hippocampal with affective information in amygdala to modulate response selection (Mogenson et al., 1980). Although our framework emphasizes the gating function of BG outputs on cortical representations, we also agree that **inputs** to the VS can be gated by contextual information from hippocampus (O'Donnell & Grace, 1995).

In previous work, we formulated a concrete hypothesis for how the BG and hippocampus interact (Atallah, Frank, & O'Reilly, 2004). We distinguish between ventral and dorsal BG areas, which receive preferentially from hippocampus and cortex, respectively (e.g., Groenewegen et al., 1987). Both BG areas are thought to play a modulatory role on motor responding, but based on the different representations encoded by their inputs. Thus, vBG modulates responding based on the conjunctive hippocampal representations that provide its input (O'Reilly & Rudy, 2001), while dBG provides modulation based on more elemental sensory representations. By this account, the hippocampus and the ventral BG work in concert, each contributing specialized functions according to the following principles: the hippocampus can rapidly bind information into conjunctive spatial representations, while the vBG can help modulate responding based on these spatial inputs, informed by prior reward-based learning history. This framework provides an alternative account to putative competitive interactions between BG and hippocampus (Poldrack et al., 1999; Poldrack et al., 2001; Packard & Knowlton, 2002). These two different views make different predictions and explanations of experimental outcomes (Atallah et al., 2004).

Chapter 4

Testing the Model: Cognitive Reinforcement Learning in Medicated and Non-medicated Parkinson's Disease

[1] Should you shout at your dog for soiling the carpet, or praise him when he does his business in the yard? Most dog trainers will tell you that the answer is both. The proverbial "carrot and stick" motivational approach refers to the use of a combination of positive and negative reinforcement: one can persuade a donkey to move either by dangling a carrot in front of it or by striking it with a stick. Both carrots and sticks are important for instilling appropriate behaviors in humans. For instance, when mulling over a decision, one considers both pros and cons of various options, which are implicitly influenced by positive and negative outcomes of similar decisions made in the past. Here we report that whether one learns more from positive or negative outcomes varies with alterations in dopamine levels caused by Parkinson's disease and the medications used to treat it.

To better understand how healthy people learn from their decisions (both good and bad), it is instructive to examine under what conditions this learning is degraded. Notably, patients with Parkinson's disease are impaired in cognitive tasks that require learning from positive and negative feedback (Knowlton et al., 1996; Ashby et al., 1998; Shohamy, Myers, Grossman, Sage, Gluck, & Poldrack, 2004). A likely source of these deficits is depleted levels of the neuromodulator dopamine in the basal ganglia of Parkinson's patients (Kish et al., 1988), because dopamine plays a key role in reinforcement learning processes in animals (Wise & Rompre, 1989). A simple prediction of this account is that cognitive performance should improve when patients take medication that elevates their dopamine levels. However, a somewhat puzzling result is that dopamine medication actually worsens performance in some cognitive tasks, despite improving it in others (Swainson et al., 2000; Cools et al., 2001).

Computational models of the basal ganglia/dopamine system provide a unified account that reconciles the above pattern of results and makes explicit predictions about the effects of medication on "carrot and stick" learning (Frank, in press; O'Reilly & Frank, 2003). These models simulate transient changes in dopamine that occur during positive and negative reinforcement, and their differential effects on two separate pathways within the basal ganglia system. Specifically, dopamine is excitatory on the direct or "Go" pathway, which helps facilitate responding, while it is inhibitory on the indirect or "NoGo" pathway, which suppresses responding (Gerfen, 1992; Aubert et al., 2000; Hernandez-Lopez et al., 1997; Hernandez-Lopez et al., 2000). In animals, phasic bursts of dopamine cell firing are observed during positive reinforcement (Schultz, 2002; Schultz et al., 1997) which are thought to act as "teaching signals" that lead to the learning of rewarding behaviors (Schultz, 2002; Wickens, 1997). Conversely, choices that no not lead to reward (and aversive events, according to some studies (Ungless, Magill, & Bolam, 2004)) are associated with dopamine dips that drop below baseline (Schultz, 2002; Satoh, Nakai, Sato, & Kimura, 2003). Similar dopamine-dependent processes have been inferred to occur in humans during positive and negative reinforcement (Holroyd & Coles, 2002; Zald, Boileau, El-Dearedy, Gunn, McGlone, Dichter, & Dagher, 2004). In our models, dopamine bursts increase synaptic plasticity in the direct pathway while decreasing it in the indirect pathway (Nishi et al., 1997; Centonze et al., 2001), supporting "Go" learning to reinforce the good choice. Dips in dopamine have the opposite effect, supporting "NoGo" learning to avoid the bad choice (Frank, in press; O'Reilly & Frank, 2003).

[1] A version of this chapter is in press in **Science** (Frank, Seeberger, & O'Reilly, 2004)

A central prediction of our models is that non-medicated Parkinson's patients are impaired at learning from positive feedback (bursts of dopamine; "carrots"), due to reduced levels of dopamine. However, the models also make the counterintuitive prediction that patients should display enhanced learning from negative feedback (dips in dopamine; "sticks"), due to their low dopamine levels that facilitate this kind of learning. Conversely, we predict that patients on medication have sufficient dopamine to learn from positive feedback, but would be relatively impaired at learning from negative feedback due to the medication blocking the effects of normal dopamine dips. This pattern of dopamine effects explains the existing puzzling results in the Parkinson's disease literature showing both cognitive enhancements and impairments from medication (Frank, in press).

This chapter presents a more direct test of the model's predictions. We employed "procedural learning" (i.e., trial and error) tasks[2] with 30 Parkinson's patients and 19 healthy seniors matched for age, education, and an estimate of verbal IQ (see Table 4.1 for demographic details and N's per task condition). Two different procedural learning tasks were used, one probabilistic and one deterministic, with the task selected at random for the first session. A subset of participants returned for a second session to perform the other task, and Parkinson's patients in this session abstained from taking their regular dose of dopamine medication for a mean of 18 hours prior to the experiment (Cools et al., 2001).

4.1 Methods

Procedures were approved by the HealthONE Institutional Review Board and the University of Colorado Human Research Committee. We tested 30 PD patients and 19 healthy seniors matched for age, education, and scores on the North American Adult Reading Test (NAART), an estimate of premorbid verbal IQ (Blair & Spreen, 1989). The demographics of seniors and PD patients are shown in Table 4.1. PD patients were recruited from the Colorado Neurological Institute. All patients were receiving daily L-Dopa preparations, with some supplemented with D2 receptor agonists and/or selegiline (monoamine activity enhancer), and were stable on their medication dose for at least 2 months. Participating senior controls were either the spouses of PD patients or were recruited from the Boulder Senior Center. Exclusionary criteria were as follows:

- significant medical history not related directly to PD (e.g. stroke, head injury, clinical dementia, epilepsy);
- concurrent illness such as schizophrenia and manic depression;
- documented or suspected history of drug abuse and/or alcoholism;
- PD patients with advanced symptoms (stage IV or V in the Hoehn and Yahr rating scale);
- PD patients with Mini Mental State Examination (MMSE) ratings of less than 24 to screen for dementia;
- patients and control subjects taking additional medication likely to confound interpretation of the findings were excluded to the best of our ability.

Participants were tested in two separate experimental sessions, separated by a minimum of 7 days. To minimize potential learning effects between sessions, we employed two different procedural learning tasks in the two sessions. The order of these tasks is randomized, so that each patient is equally likely to perform either of the tasks ON or OFF medication, and each healthy senior participant is equally likely to perform each task in the first or second session. Both tasks require trial and error learning, in which some stimuli have a net positive reinforcement value (and should be chosen), whereas others have a negative reinforcement value (and should

[2] Note that although participants may become aware of which stimulus they are actually choosing, they are likely not aware of the detailed influence of the history of positive and negative feedback on their choices. This is the sense in which we regard these as procedural or implicit learning tasks.

Group/Task	n	n filt	Sex ratio (m:f)	Age	Years Education	NAART IQ	Hoehn & Yahr ratio (1:2:2.5)
Seniors	19		10:9	64.8 (1.3)	17.1 (0.6)	118.5 (1.6)	N/A
PS	15	11	5:6	64.9 (1.1)	16.7 (0.9)	119.0 (1.9)	N/A
TI	18	17	9:8	64.1 (1.4)	16.7 (0.6)	118.4 (1.7)	N/A
PD patients	30		19:11	63.8 (2.2)	16.9 (0.4)	114.5 (1.9)	1:22:7
ON	29*	26	17:9	62.6 (2.6)	16.9 (0.5)	112.6 (2.0)	1:18:7
PS	12	9	7:2	69.1 (2.6)	16.9 (0.8)	112.3 (3.4)	0:6:3
TI	17	17	10:7	59.2 (3.5)	16.8 (0.7)	112.7 (2.6)	1:12:4
OFF	17	17	10:7	65.6 (3.3)	17.3 (0.7)	115.9 (2.5)	0:14:3
PS	9	9	5:4	62.3 (5.2)	16.8 (0.9)	113.2 (3.9)	0:8:1
TI	8	8	6:2	70.5 (1.5)	17.4 (0.9)	119.7 (1.9)	0:6:2

Table 4.1: Demographic variables for seniors and PD patients, with no significant differences between groups in any of the demographic variables. Imbalance in N's resulted from random assignment to task in session 1, together with differences in filtering. Because only a subset of participants returned for the second session, variables are broken down to show the number of participants in each condition that actually performed each task (PS = probabilistic selection; TI = transitive inference). The "n filt" column shows the number of remaining participants who were not filtered out for data analysis (see Data Filtering sections); participants who were filtered out were not included in the demographic means displayed for that row (as they were not used in the statistical comparisons). NAART IQ = premorbid IQ as estimated with the North American Adult Reading Test. For PD patients, disease severity is indicated in terms of the number of patients classified in Hoehn and Yahr stages 1, 2, and 2.5. Values represent mean (standard error). * Although other patients participated ON medication in the first session, one patient was OFF medication on that day (and did not return ON medication), explaining the discrepancy between the N for total PD patients and the N for patients ON medication. Because of a lack of session effect in either controls or patients, this does not confound any of our reported results.

be avoided). Both tasks use two-alternative forced choice, in which participants press one of two keys on a keyboard to "choose" one of two stimuli presented on a computer screen.

Participants sit in front of a computer screen in a lighted room and view pairs of visual stimuli that are not easily verbalized (Japanese Hiragana characters. These stimuli are presented in black on a white background, in 72 pt font. They press keys on the left or right side of the keyboard depending on which stimulus they choose to be "correct". Note that precise motor control is not necessary because any of 12 keys on the appropriate half of the keyboard counts as a response, allowing us to control for motor deficits associated with PD. Furthermore, the forced-choice nature of the task controls for any differences in overall motor responding. Visual feedback is provided following each choice (the word "Correct!" printed in blue or "Incorrect" printed in red). If no response is made within four seconds, the words "no response detected" are printed in red.

4.1.1 Task I: Probabilistic Selection

4.1.1.1 Rationale

In previous work, we explored the involvement of BG/DA interactions in probabilistic classification via explicit computational modeling, which (a) accounted for PD deficits and (b) made novel predictions (Frank, in press). Specifically, the model predicted different deficits in medicated and non-medicated PD patients, depending on whether the learning was from positive or negative feedback. Due to depleted DA, those OFF medication should have deficits in learning from positive feedback. DA medication should alleviate these deficits, but may impair learning from negative feedback as the medication blocks the effects of DA dips that are required for this learning (Frank, in press). [3] We therefore designed a novel "Probabilistic Selection" (PS) task to specifically test this hypothesis.

4.1.1.2 Procedures

In the Probabilistic Selection task, three different stimulus pairs (AB, CD, EF) are presented in random order and participants have to learn to choose one of the two stimuli (Figure 4.1a). Feedback follows the choice to indicate whether it was correct or incorrect, but this feedback is probabilistic. In AB trials, a choice of stimulus A leads to correct (positive) feedback in 80% of AB trials, whereas a B choice leads to incorrect (negative) feedback in these trials (and vice-versa for the remaining 20% of trials). CD and EF pairs are less reliable: stimulus C is correct in 70% of CD trials, while E is correct in 60% of EF trials. Over the course of training participants learn to choose stimuli A, C and E more often than B, D, or F. Note that learning to choose A over B could be accomplished either by learning that choosing A leads to positive feedback, or that choosing B leads to negative feedback (or both). To evaluate whether participants learned more about positive or negative outcomes of their decisions, we subsequently tested them with novel combinations of stimulus pairs involving either an A (AC, AD, AE, AF) or a B (BC, BD, BE, BF); no feedback was provided. We predict that Parkinson's patients on medication learn more from positive feedback, and should therefore reliably choose the best "carrot" (stimulus A) in all novel test pairs in which it is present. In contrast, those off medication learn more from negative feedback, and should therefore reliably avoid the worst "stick" (stimulus B).

We enforced a performance criterion (evaluated after each training block of 60 trials) to ensure that all participants were at the same performance level before advancing to test. Because of the different probabilistic structure of each stimulus pair, we used a different criterion for each (65% A in AB, 60% C in CD, 50% E in EF)[4] .

[3] Note that although aversive events have been reported to increase firing in a small proportion of cells in DA-producing brain areas (Schultz, 2002), these neurons are in fact non-dopaminergic; DA cells in this region are indeed inhibited during aversive events (Ungless et al., 2004). Further, we argue that "Incorrect" feedback signals in trial-and error tasks are not actively aversive, and may be more analogous to a lack of positive reinforcement, widely accepted to induce DA dips in animal studies (Schultz, 2002; Satoh et al., 2003).

[4] In the EF pair, stimulus E is correct 60% of the time, but this is particularly difficult to learn. We therefore used a 50% criterion

After reaching this criterion, participants were subsequently tested with the same training pairs, in addition to all novel combinations of stimuli, in random sequence. They were instructed (prior to the test phase) to use "gut instinct" if they did not know how to respond to these novel pairs. Each test pair was presented 6 times.

4.1.2 Task II: Transitive Inference

4.1.2.1 Rationale

Our framework holds that PD patients are not impaired at learning probabilistic tasks **per se**. Rather, we argue that these tasks are simply examples of situations in which choices for particular stimuli are only differentiated by subtle reinforcement histories. We therefore predicted that other non-probabilistic tasks that engender subtle differences in associative strength would also be learned in a qualitatively different way by PD patients.

In the Implicit Transitive Inference task (Frank et al., in press), the reinforcement for each stimulus pair is deterministic, but stimulus pairs are partially overlapping Figure 4.1a. Four pairs of stimuli are presented: A+B-, B+C-, C+D- and D+E- where + and – indicate positive and negative feedback. A hierarchy (A>B>C>D>E) emerges in which stimuli near the top of the hierarchy develop net positive associative strengths, while those near the bottom develop net negative associative strengths (Frank et al., in press; Frank et al., 2003; von Fersen et al., 1991). This explains why, when presented with novel combination BD, participants often correctly choose stimulus B, despite having no explicit awareness of any hierarchical structure among the items (Frank et al., in press; Frank et al., 2003; Wynne, 1995; Delius & Siemann, 1998). Indeed, the same associative principles may give rise to transitive responding in animals trained in similar paradigms, and computational models (VanElzakker, O'Reilly, & Rudy, 2003; Frank et al., 2003; von Fersen et al., 1991; Wynne, 1995, 1998; Delius & Siemann, 1998; Siemann & Delius, 1998). Based on this associative account, we predicted that Parkinson's patients on medication would learn more about the positive associations at the top of the hierarchy, resulting in better performance on stimulus pairs AB and BC. Conversely, those off medication should learn more about the negative associations at the bottom of the hierarchy, resulting in better CD and DE performance. Note that because the novel BD pair could be solved either by a positive B association or a negative D association, we did not predict a difference in BD performance between groups.

4.1.2.2 Procedures

Training consisted of four phases of blocked trials, followed by a fifth phase of randomly interleaved trials. Each phase was terminated after criterion performance of at least 75% correct across all pairs was achieved. In the first phase, the stimulus pairs were presented in blocks of 6 trials, such that the first block consisted of AB trials, the second block consisted of BC trials, and so on. In phase 2, blocks were shortened to 4 trials per block. Phase 3 consisted of 3 trials per block, and phase 4 consisted of 2 trials per block. In phase 5, all pairs were randomly interleaved for a total of 25 trials before criterion performance was evaluated. If criterion was not met, the random sequence was repeated. The test phase was similar to the training phase 5 in that all pairs were randomly interleaved. However, no feedback was provided and the two transitive pairs BD and AE were added to the mix of randomly ordered pairs. All pairs were presented 6 times each. Following the experiment, all participants were given a questionnaire to assess their awareness of the logical hierarchy of the stimuli, and to determine what strategies, if any, were used to respond to the novel test pairs (Frank et al., in press, Appendix B). Out of 17 seniors, only one became explicitly aware of the hierarchical structure. No PD patients became explicitly aware.

---

for this pair simply to ensure that if participants happened to "like" stimulus F at the outset, they nevertheless had to learn that this bias was not going to consistently work.

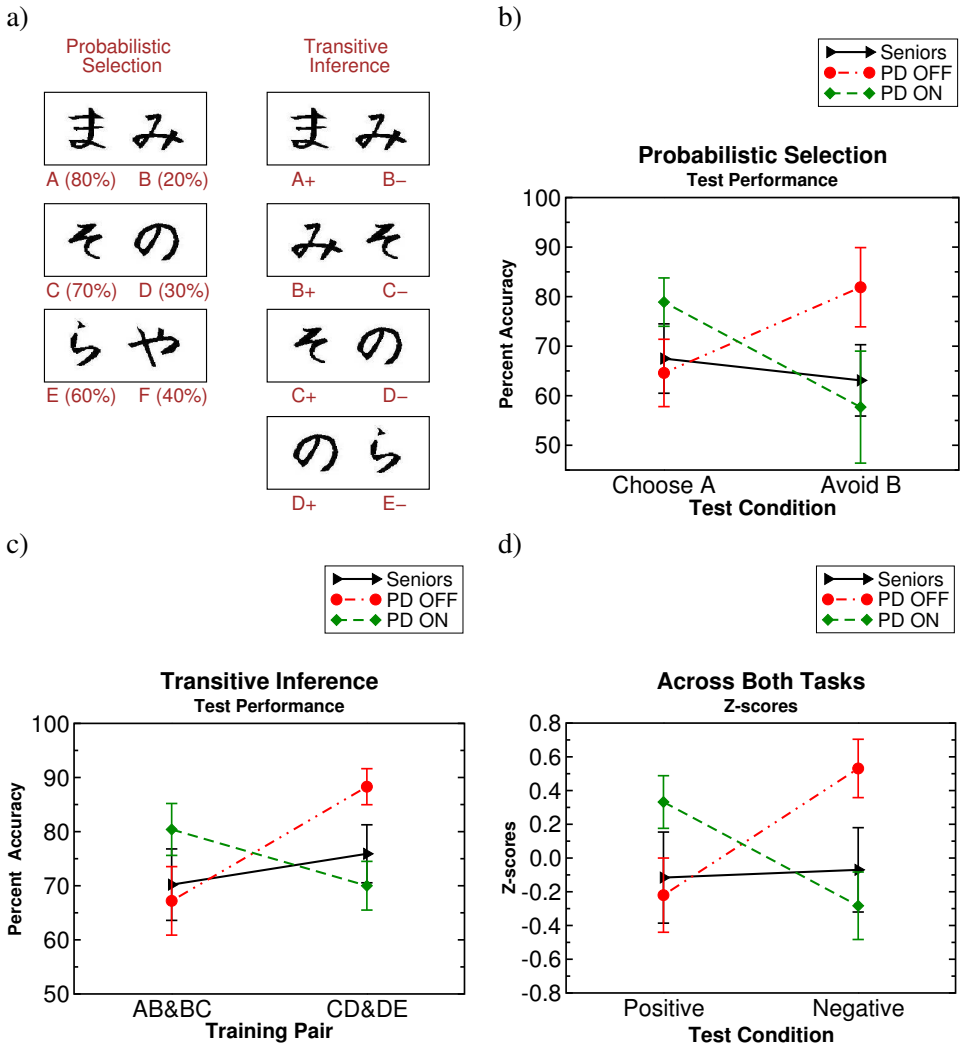

Figure 4.1: **a)** Example stimulus pairs (Hiragana characters) used in both cognitive procedural learning tasks, designed to minimize verbal encoding. One pair is presented per trial, and the participant makes a forced choice. In Probabilistic Selection, the frequency of positive feedback for each choice is shown. In Transitive Inference, feedback is deterministic and indicated by the +/- signs for each stimulus. Any of 12 keys on the left side of the keyboard selects the stimulus on the left, and vice-versa for the right stimulus. The stimulus locations were randomized across trials, and assignment of Hiragana character to stimulus label (A-F) was randomized across participants. In actuality, different Hiragana characters were used across tasks. **b)** Novel test pair performance in the probabilistic selection task, where choosing A depends on having learned from positive feedback, while avoiding B depends on having learned from negative feedback. **c)** Training pair performance during the test phase in the transitive inference task. Stimuli at the top of the hierarchy (A, B) have net positive associations, while those at the bottom (C, D) have net negative associations (Frank et al., in press; Frank et al., 2003; von Fersen et al., 1991; Wynne, 1995; Delius & Siemann, 1998). Thus, learning from positive feedback benefits performance on AB & BC, while learning from negative feedback benefits CD & DE. Groups did not differ in novel test pairs AE and BD (not shown; see Table 4.3) which could be solved either by choosing stimuli with positive associations or avoiding those with negative associations. **d)** Z-scores across both probabilistic selection and transitive inference tasks. Positive and negative conditions correspond to A & B pairs in the probabilistic selection task, and AB/BC & CD/DE pairs in the transitive inference task. Error bars reflect standard error.

4.2 Main Results

Results confirmed our predictions. Despite no main effect of medication, session, or test condition, the critical interaction between medication and test condition was significant for both the probabilistic selection [$F(1,26) = 4.3$, $p < .05$] and transitive inference [$F(1,39) = 5.5$, $p < .05$] tasks. In the probabilistic selection task (Figure 4.1b), patients on medication chose stimulus A no matter what it was paired with, indicating that they had found the best carrot in the bunch. In contrast, patients off medication had a greater tendency to avoid stimulus B no matter what it was paired with, indicating that they had learned to avoid the harshest stick. Aged matched controls did not differ in performance between A and B pairs. In the transitive inference task (Figure 4.1c), patients on medication performed better at choosing positively associated stimuli at the upper end of the hierarchy, whereas those off medication more reliably avoided negative stimuli at the lower end. Finally, aged matched controls did not differ between performance on pairs at the high and low end of the stimulus hierarchy. There was also no effect of medication on performance on novel pairs AE and BD [$F(1,39) = 1.6$, n.s.].

To compare results across both tasks, we converted accuracy measures for both positive and negative conditions to Z-scores (Figure 4.1d). This analysis confirmed a significant interaction between positive/negative condition and Parkinson's disease medication group [$F(1,68) = 10.4$, $p = 0.0019$]. Planned comparisons revealed that patients on medication chose positive stimuli more reliably than they avoided negative stimuli [$F(1,25) = 4.98$, $p < .05$], and more reliably than the other two groups [$F(1,69) = 4.8$, $p < .05$]. Conversely, patients off medication avoided negative stimuli more reliably than they chose positive stimuli [$F(1,15) = 5.42$, $p < .05$], and more reliably than the other two groups [$F(1,69) = 7.6$, $p < 0.05$].

This last observation is a rare example of enhanced cognitive performance associated with neurological disease, as it suggests that non-medicated patients made better use of negative feedback. Trial-to-trial analysis confirmed that a change of choice behavior in the probabilistic selection task (e.g., they chose C in a CD trial after having chosen D in the previous CD trial), was more often accounted for by negative feedback in the previous trial in patients off medication compared with those on medication [$F(1,26) = 5.62$, $p < .05$]. Medicated patients switched choices just as often during training, but were not as influenced by negative feedback to do so. There was no difference between these groups in the efficacy of positive feedback to modify behavior on a trial-to-trial basis [$F(1,26) = 0.42$, n.s.].

4.2.1 Detailed Results and Analysis: Probabilistic Selection

Reported below are detailed results and analysis used to provide the basis for the main results reported above.

4.2.1.1 Training

During the training phase of the PS task, all PD patients and seniors successfully learned to choose positively associated stimuli more often than negatively associated stimuli. Groups did not significantly differ in the number of training trials needed to reach criterion [OFF: 200(40); ON: 273(33); SEN: 250(48), $F(2,26) = 0.75$, n.s.]. In this and all subsequent cases, numbers in parentheses are standard errors (SEM). There was also no between-groups differences in accuracy, either across all training conditions [$F(2,26) = 0.55$, n.s.], or when considering just the critical AB pair alone [$F(2,26) = 0.74$, n.s.].

4.2.1.2 Data Filtering

Because we were interested in the extent to which participants learned about the positive versus negative outcomes of their choices, we had to first ensure that they learned the basic task. While the training criteria were meant to address this issue, some participants were globally confused by the lack of feedback and addition

Group	Session 1	Session 2
Seniors	0.021 (0.21)	-0.29 (0.2)
PD	0.026 (0.11)	0.15 (0.12)

Table 4.2: Session effects on test phase accuracy, converted to Z-scores to permit comparison across tasks. Values represent mean (standard error).

of novel pairs during test and therefore performed poorly all around, including in pairs which were easiest for them during training. To reduce the amount of noise caused by this confound, we eliminated participants from the analysis who did not perform better than chance during test in the easiest training pair conditions. In the PS task, we eliminated three patients ON medication and four seniors who did not choose A over B more than 50% of the time when the AB pair was presented at test, reasoning that if they could not reliably choose A/avoid B in this pair, then the results in novel pairs were meaningless. The results described below apply to the remainder of participants, amounting to 11 seniors, 9 patients ON medication and 9 patients OFF medication.

4.2.1.3 Session Effects

To minimize practice and learning effects across session, each participant performed one task in the first session and the other task in the second. However, it is nevertheless possible that there were non-specific transfer effects across sessions that are unrelated to the particular details of each task. To examine this possibility, we performed an ANOVA across all participants to test for a main effect of session on overall test performance. We found no such effect [$F(1,27) = 0.10$, n.s.]. As shown in Table 2, the general trend for better performance of OFF medication patients relative to ON patients is evident by greater Z-scores in the second session. However, senior participants showed the reverse (non-significant) trend, supporting our claim that performance differences in the patient group were related to medication effects, rather than generalized transfer effect that theoretically could have carried over from the first session. Because patients in the second session were always OFF medication, we did not include session number in subsequent analyses, as it would likely mask any legitimate medication effects.

4.2.1.4 Test Pair Analysis

Next, we performed a general linear model (GLM) regression on positive/negative test pair accuracy, using between-subjects contrast-coded factors of patients versus seniors and patients ON versus OFF medication (forming a full set of orthogonal contrast codes) (Judd & McClelland, 1989). There were no overall differences in performance levels between patients and seniors [$F(1,26) = 1.0$, n.s.], or between patients ON and OFF medication [$F(1,26) = 0.57$, n.s.]. There was no within-subjects main effect of positive/negative (A/B) test condition [$F(1,26) = 0.14$, n.s.]. Critically, within-subjects differences in positive/negative test condition interacted significantly with the ON/OFF medication status [$F(1,26) = 4.33$, $p = .04$] (see Figure 1b in the main paper). There was a trend for patients ON medication to perform better than those OFF medication at choosing positive stimuli [$F(1,26) = 2.3$, $p = 0.14$], while OFF patients performed better than ON patients when having to avoid negative stimuli [$F(1,26] = 3.5$, $p = .07$]. Finally, to explicitly compare performance of each of the medication conditions with the senior control participants, we performed two additional GLM's to test for the specific contrasts of patients ON medication versus seniors and patients OFF medication versus seniors (each of these tests was also accompanied by the orthogonal contrast code). Patients ON medication were non-significantly better than healthy seniors at choosing positive stimuli [$F(1,26) = 1.6$, n.s.] but worse than seniors at avoiding negative stimuli [$F[1,26] = 0.2$]. In contrast, patients OFF medication were numerically worse than seniors at choosing positive stimuli [$F(1,26) = 0.11$], but better than seniors at avoiding negative stimuli [$F(1,26) = 2.1$, $p = 0.16$].

4.2.1.5 Trial-to-trial Go/NoGo Learning from Positive/Negative Feedback

If our hypothesis is correct, we could potentially find evidence of it on a trial-to-trial basis during training, in terms of the effectiveness of positive/negative feedback in one trial to modify behavior in the subsequent trial of the same type. In other words, how much of participants' choice behavior in a given trial could be predicted by the type of feedback in the previous trial of the same type?

We restricted our analysis to the first 100 trials of training, because participants are less sensitive to feedback from individual trials in the latter parts of training as choice behavior stabilizes. We analyzed the conditional probabilities that (a) feedback on the previous trial was positive given that participants chose the same stimulus in the current trial; (b) feedback on the previous trial was negative given that participants switched to the alternative stimulus on the current trial. For the positive feedback case (a), the conditional probabilities were SEN: 55.4 (4.36); PD OFF 57.7 (3.7); PD ON 54.1 (2.3). There were no significant differences between PD groups in the efficacies of positive feedback in modifying behavior [$F(1,26) = 0.42$, n.s.]. For the negative feedback case (b), the conditional probabilities were SEN: 58.2 (2.44); PD OFF 62.0 (4.15); PD ON 49.4 (4.41). This difference between patients OFF and ON medication was significant [$F(1,26) = 5.62$, $p = .025$].

The enhancement in learning from positive feedback for the PD ON group was not apparent on a trial-to-trial basis during training, but was evident during test choice behavior. This suggests that greater levels of DA in medicated patients were effective when positive feedback had to be integrated over several trials, but were not sufficient to modify behavior after a single trial. In contrast, non-medicated patients were significantly more likely to avoid a negative stimulus, and this was evident both during test and in the effectiveness of a single negative feedback signal to modify behavior during training. That trial-to-trial effects were found for negative, but not positive, feedback is consistent with the fact that the tendency to avoid negative stimuli by OFF patients was generally greater than the tendency to choose positive stimuli by ON patients.

4.2.2 Detailed Results and Analysis: Transitive Inference

4.2.2.1 Training

During the training phase of the TI task, both Parkinson patients and aged-matched controls successfully learned to choose correctly across all stimulus pairs. Again, groups did not significantly differ in the number of training trials needed to reach the performance criteria required to advance to test [OFF: 280 (44); ON: 359 (35); SEN: 302 (39), $F(2,39) = 1.03$, n.s.]. There was also no between-groups differences in training phase accuracy [$F(2,39) = 0.55$, n.s.].

4.2.2.2 Data Filtering

As in the PS task, we filtered out participants who were globally confused by the test phase and did not perform better than chance even at the easiest training pair conditions during test. This amounted to filtering out only one senior who did not perform better than 50% measured across anchor pairs AB and DE.[5] Results described below apply to the remainder of participants, amounting to 17 seniors, 17 patients ON medication and 8 patients OFF medication.

4.2.2.3 Session Effects

As in the PS task, an initial simple ANOVA revealed no main effects of session number [$F(1,40) = 0.02$, n.s.] on overall test pair performance.

[5] The reason only one participant failed to transfer training knowledge to test performance in the TI task, while a total of seven failed in the PS task, may relate to the addition of novel test pairs. In the PS task, eight of the test pairs were novel with only three pairs that repeated from the training phase. In the TI task, there were only six test pairs, of which four of them were the original training pairs.

Group	AE	BD
Seniors	98 (1.3)	61.8 (10.0)
PD OFF	95.8 (2.7)	66.2 (12.9)
PD ON	84.3 (5.8)	58.4 (9.5)

Table 4.3: Novel test pair results in the transitive inference task. Values represent mean (standard error). As predicted, groups did not differ in test pair generalization, since both AE and BD could be solved either by choosing the positive stimulus or by avoiding the negative stimulus.

4.2.2.4 Test Pair Analysis

As in the PS task, we performed a GLM regression on positive/negative test pair accuracy, testing the specific contrasts of patients versus seniors, and patients ON versus OFF medication. There were no overall differences in performance levels between patients and seniors [F(1,39) = 0.62, n.s.], or between patients ON and OFF medication [F[1,39] = 0.13, n.s.]. There was no within-subjects main effect of positive/negative (A/B) test condition [F(1,39) = 1.51, n.s.]. Critically, within-subjects differences in positive/negative test condition interacted significantly with the ON/OFF medication status [F(1,39) = 5.54, p = .023] (see Figure 1c in the main paper). Patients ON medication performed non-significantly better than those OFF medication at choosing positive stimuli [F(1,39) = 1.6, p = 0.2], while OFF patients performed significantly better than ON patients when having to avoid negative stimuli [F(1,39] = 4.2, p = .047]. Finally, in comparison with healthy seniors, patients ON medication were non-significantly better than healthy seniors at choosing positive stimuli [F(1,39) = 1.6, n.s.] but worse than seniors at avoiding negative stimuli [F[1,39] = 0.5]. In contrast, patients OFF medication were numerically worse than seniors at choosing positive stimuli [F(1,39) = 0.1], but better than seniors at avoiding negative stimuli [F(1,26) = 2.3, p = 0.14].

The results for novel test pairs BD and AE are shown in Table 4.3. There was no significant effect of PD group on test pair performance [F(1,39) = 1.61, n.s.]. Note that we did not predict a difference in test pair performance, because participants can choose B over D or A over E **either** by learning that A and B are good **or** by learning that D and E are bad.

4.2.2.5 Relation to Model Predictions

Further insight comes from analysis of the more detailed pattern of results for each training pair (Figure 4.2), which was predicted by our computational models. The intact model learned "Go" to the stimulus at the top of the hierarchy (A) and NoGo to the one at the bottom (E). This enabled the model to allow B to take on a net positive value to facilitate performance in the BC pair, because this positive B value would not interfere with the very strong positive strength associated with A. A similar process held for the D item in the CD case, which took on a net negative value that did not compete with the very strong negative value for E. The resulting positive B and negative D values explains why B is chosen over D in implicit versions of the TI task in humans (Frank et al., in press) and prior studies with animals (Delius & Siemann, 1998; Frank et al., 2003; VanElzakker et al., 2003; von Fersen et al., 1991).

In contrast to this intact case, models with simulated PD (OFF medication) were biased to learn with a "NoGo strategy", such that they chose A in AB indirectly by learning NoGo to B. This led to worse performance on the BC pair, in which stimulus B should be chosen. In contrast, this NoGo bias helped performance on the CD pair by increasing the model's tendency to avoid D. This is exactly the pattern seen in Figure 4.2, where the OFF medication PD patients perform worse at BC, and better at CD.

Conversely, models with simulated DA medication were biased to learn with a "Go strategy", such that they learned Go to D in the DE pair instead of NoGo to E. This led to worse performance on the CD pair. However, the Go bias improved performance in the BC case, leading to a stronger tendency to select B. Again, this is the pattern seen in the ON medication PD patients in Figure 4.2, where they perform better at BC than

CD, reversing the pattern seen in the OFF medication group.

4.2.3 Detailed Results and Analysis: Combined Z-scores

Because we wanted to examine the extent to which participants learned from positive versus negative feedback across both tasks, we needed a way to equate performance on positively and negatively associated cues in the two tasks. To do this, we first computed Z-scores for performance in the two conditions for each task alone, by normalizing accuracy measures so that the distribution across all participants had a mean of 0 and a standard deviation of 1. In the PS task, we computed participants' Z-scores for test performance on A pairs and B pairs, while in the TI task we computed Z-scores for test performance on AB/BC pairs and CD/DE pairs. We then combined Z scores across the two tasks in positive and negative conditions, which were now in the same metric (Figure 1d in the main paper).

We first performed a GLM regression to test for within-subject differences in positive versus negative test pair accuracy, and whether these differences interacted with group (patients versus seniors) or medication status (ON versus OFF). Across all participants, there was no within-subject main effect of positive/negative test condition [$F(1,68) = 0.3$, n.s.]. Critically, there was a highly significant interaction between Z-scores on positive/negative test conditions and PD medication status (ON/OFF) [$F(1,68) = 10.4$, $p = 0.0019$]. Each of the simple main effects were also significant when analyzed across both tasks: planned contrasts revealed that patients ON medication had significantly greater Z-scores for choosing positive stimuli than those OFF medication [$F(1,68) = 4.3$, $p = .04$]. Conversely, when avoiding negative stimuli, OFF patients had significantly greater Z-scores than ON patients [$F(1,68) = 8.0$, $p = .006$].

For further between-groups analysis, we performed two additional planned pair-wise comparisons across all participants—one for accuracy on positive test pairs, and one for accuracy on negative test pairs. For positive pairs, we tested whether Z-scores of patients ON medication were better than those of the other two groups. For negative pairs, we tested whether Z-scores of patients OFF medication were better than those of the other two groups. These analyses revealed that patients ON medication had significantly greater positive Z-scores than the other two groups combined [$F(1,69) = 4.8$, $p = .03$], while those OFF medications had significantly greater negative Z-scores than the other two groups [$F(1,69) = 7.6$, $p = .007$]. Compared with just the senior group, patients ON medication had marginally higher positive Z-scores [$F(1,69) = 3.3$, $p = .07$] and non-significantly lower negative Z-scores [$F(1,69) = 0.7$]. Conversely, when compared with seniors, those OFF medication had significantly greater Z-scores for avoiding negative stimuli [$F(1,69) = 4.6$, $p = .035$], with numerically but non-significantly lower positive Z-scores [$F(1,69) = 0.2$].

Finally, we also performed separate repeated measures ANOVA's for Z-scores of each medication condition, to evaluate whether patients ON medication were better at choosing positive stimuli than they were at avoiding negative stimuli (and vice versa for patients OFF medication). Indeed, across both tasks, patients ON medication were better at choosing positive stimuli than these same patients were at avoiding negative stimuli [$F(1,25) = 4.98$, $p = .03$]. Conversely, those OFF medication better avoided negative stimuli than they chose positive stimuli [$F(1,15) = 5.42$, $p = .03$]. Healthy seniors did not differ in their performance on positive versus negative stimuli [$F[1,18] = 0.02$, n.s.].

4.3 Discussion

Taken together, these findings provide a mechanistic understanding of the nature of the cognitive sequelae of Parkinson's disease, which ties together a variety of other observations across multiple levels of analysis. First, we build on claims that learning from error feedback is primarily affected in Parkinson's disease (Shohamy et al., 2004), by showing that the direction of this effect interacts critically with the valence of the feedback and the medication status of the patient. Second, these results are consistent with neuroimaging studies showing that positive and negative feedback have differential effects on basal ganglia activity (Delgado, Locke, Stenger, & Fiez, 2003; Zink, Pagnoni, Martin-Skurski, C., & Berns, 2004). Third, they help clarify the

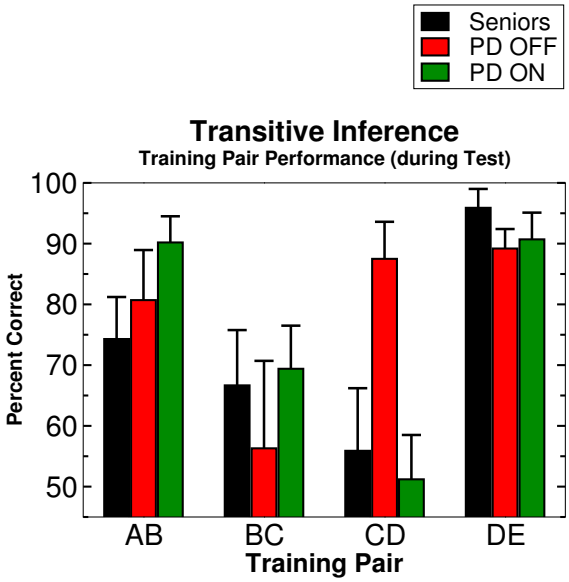

Figure 4.2: Training pair performance during the test phase of the transitive inference task. Patterns were consistent with model predictions: while all participants performed relatively well at the anchor pairs AB and DE, patients OFF medication performed poorly on the BC pair, whereas those ON medication had difficulty on the CD pair. See text for explanation.

basis for why medication sometimes improves but sometimes impairs cognitive deficits in Parkinson's disease, depending on the task (Swainson et al., 2000; Cools et al., 2001; Frank, in press). Specifically, patients on medication displayed enhanced positive feedback learning beyond even that of healthy participants, supporting the idea that medication results in higher than normal amounts of dopamine in ventral striatum, which is relatively spared in early stage Parkinson's disease (Kish et al., 1988; Swainson et al., 2000; Cools et al., 2001). Finally, our observation that non-medicated patients display enhanced ability to avoid negative stimuli may provide the fundamental basis for reports of enhanced harm avoidance behavior in these patients (Kaasinen, Nurmi, Bergman, Eskola, Solin, Sonninen, & Rinne, 2001; Tomer & Aharon-Peretz, 2004).

An equally important contribution of this work is in its confirmation of very specific predictions made by our computational model of the basal ganglia system (Frank, in press; O'Reilly & Frank, 2003; Frank et al., 2001). Almost all of the basic mechanisms of this model have been postulated in various forms by other researchers. Nevertheless, it represents an integration of these mechanisms into a coherent, mechanistically explicit system. At the most general level, the basal ganglia in our model modulates the selection of actions being considered in frontal cortex (Mink, 1996; Chevalier & Deniau, 1990; Frank et al., 2001; Frank, in press). More specifically, two main projection pathways from the striatum go through different basal ganglia output structures on the way to thalamus and up to cortex (Figure 4.3a). Activity in the direct pathway sends a "Go" signal to facilitate the execution of a response considered in cortex, whereas activity in the indirect pathway sends a "NoGo" signal to suppress competing responses. Transient changes in dopamine levels that occur during positive and negative feedback have opposite effects on D1 and D2 receptors, which are relatively segregated in the direct and indirect pathways, respectively (Gerfen, 1992; Aubert et al., 2000; Hernandez-Lopez et al., 1997; Hernandez-Lopez et al., 2000). Thus the net effect of dopamine bursts during positive reinforcement are to activate the Go pathway and deactivate the NoGo pathway, driving learning so that reinforced responses are subsequently facilitated. Conversely, decreases in dopamine during negative reinforcement have the opposite effect, driving NoGo learning so that incorrect responses are subsequently suppressed/avoided (Frank, in press).

These dopamine modulation effects on the Go and NoGo pathways lead directly to the predictions that we confirmed in the experiments reported earlier, as revealed in computational simulations of these dynamics (Figure 4.3; Frank, in press). To simulate Parkinson's disease, we decreased tonic and phasic dopamine levels in the SNc layer of the network, which reduced the ability to generate dopamine bursts during positive feedback. Therefore, the model was relatively impaired at reinforcing Go firing to correct responses. Furthermore, the low tonic dopamine levels produced a persistent bias on the system in favor of the NoGo pathway, which resulted in a corresponding bias to learn NoGo in response to negative feedback. Thus, in our simulation of the probabilistic selection task (Figure 4.3b), the simulated Parkinson's model learned more to NoGo to B instead of Go to A (see Appendix B for statistics). In contrast, intact models learned an even balance of Go to A and NoGo to B.

To simulate the effects of Parkinson's disease medication, we increased the dopamine levels (both tonic and phasic), but we also decreased the size of the phasic dopamine dips during negative feedback. This latter effect is included because D2 agonist medications taken by the vast majority of our Parkinson's patients (in addition to L-Dopa) tonically bind to D2 receptors irrespective of phasic changes in dopamine firing, thereby "filling in" the dips. The net result is the opposite of our simulated Parkinson's model. The tonic elevation in dopamine receptor activation produced a Go bias in learning, while the diminished phasic dip decreased the model's ability to learn NoGo from negative feedback. These combined effects produced the clear crossover-interaction pattern that we observed in our studies (Figure 4.3b, see Appendix B for statistics). Similar results held for our simulation of the transitive inference experiment (see section 4.2.2.5). Finally, reversal learning deficits observed in Parkinson's patients on medication (Swainson et al., 2000; Cools et al., 2001) were also accounted for by this same model (Frank, in press).

Nevertheless, the model does not capture the overall better performance of Parkinson's patients in our study relative to healthy senior controls. This result is somewhat surprising, given that patient impairments have been observed in previous studies (Knowlton et al., 1996; Ashby et al., 1998; Shohamy et al., 2004). One potential reason for this discrepancy is the relative simplicity of our task relative to those used in previous studies.

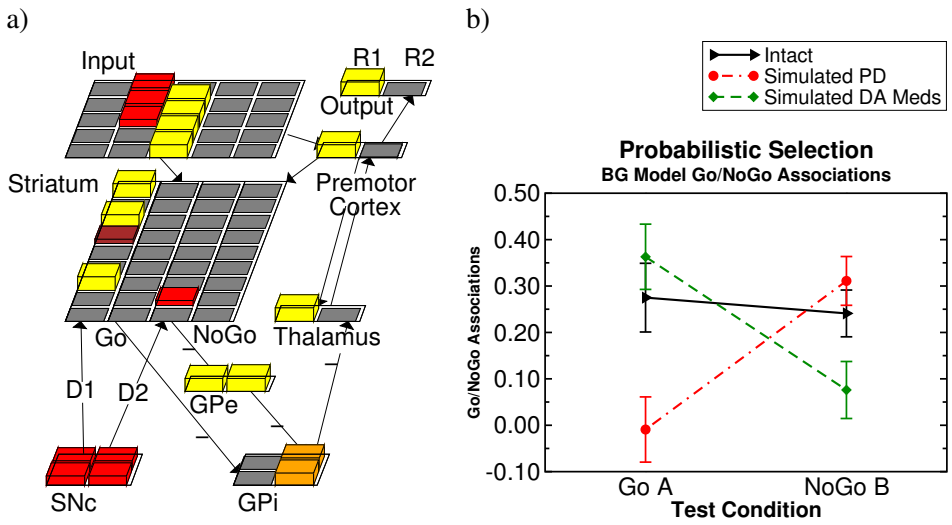

Figure 4.3: **a)** The Frank (in press) neural network model of the basal ganglia circuit (squares represent units, with height and color reflecting neural activity; yellow = most active, red = less active, grey = not active). The Premotor Cortex selects an Output response via direct projections from the sensory Input, and is modulated by the BG projections from Thalamus. Go units are in the left half of the Striatum layer; NoGo in the right half, with separate columns for the two responses (R1 (left button), R2 (right button)). In the case shown, striatum Go is stronger than NoGo for R1, inhibiting GPi, disinhibiting Thalamus, and facilitating execution of the response in cortex. A tonic level of dopamine is shown in SNc; a burst or dip ensues in a subsequent error feedback phase (not shown), causing corresponding changes in Go/NoGo unit activations, which drive learning. **b)** Predictions from the model for the probabilistic selection task, showing Go - NoGo associations for stimulus A and NoGo - Go associations for stimulus B. Error bars reflect standard error across 25 runs of the model with random initial weights.

Furthermore, although our control group was matched to the patients in all of our demographic variables, other uncontrolled variables might have led to differences in overall performance levels. For example, because we had access to patient medical records, we may have successfully excluded more patients than seniors for other age-related neurological impairments. Alternatively, Parkinson's patients may have had greater motivation to perform well, given that they were aware that we were studying cognitive sequelae of their disease (the so-called "Hawthorne Effect"). Further, while abstract neural models can make qualitative predictions (such as the cross-over interactions observed in this study), the quantitative aspects of the predictions require more detailed knowledge of specific parameters of the neural system, along with the precise degree of dopamine depletion and remediation by medication in Parkinson's disease; these data are not available. Therefore, we argue that the most meaningful comparisons are the on-vs-off medication conditions, for which the model and data are in close agreement. In addition, the model accurately predicts that healthy seniors did not differ in their tendency to learn from positive versus negative feedback. Finally, we note that our model does not explicitly consider the uneven levels of dopamine depletion in ventral and dorsal striatum of Parkinson's patients (Kish et al., 1988), which are also thought to play a role in cognitive enhancements/impairments resulting from medication (Swainson et al., 2000; Cools et al., 2001; Frank, in press).

In summary, we have presented evidence for a mechanistic account of how the human brain implicitly learns to make choices that lead to good outcomes, while avoiding those that lead to bad outcomes. The consistent results across tasks (one probabilistic and the other deterministic), and in both medicated and non-medicated Parkinson's patients, provide substantial support for a dynamic dopamine model of cognitive reinforcement learning.

Chapter 5

Testing the Model: Cognitive Reinforcement Learning and Executive Function in Healthy Participants Under Pharmacological Challenge

[1] As described in the previous chapters, the basal ganglia (BG) are thought to participate in various aspects of cognition and behavior by interacting with and modulating multiple areas of frontal cortex (Alexander et al., 1986). Similarly, the neurotransmitter dopamine (DA) plays a modulatory role in cognition through extensive diffuse projections from midbrain DA nuclei to the BG and frontal cortical areas (Joel & Weiner, 2000). Several neurological conditions implicate DA dysfunction, including Parkinson's disease (PD), ADHD, and schizophrenia (Nieoullon, 2002; Kish et al., 1988; Dougherty et al., 1999; Ilgin et al., 2001; Seeman, 1987; Weinberger, 1987; Abi-Dargham et al., 2000; McGowan et al., 2004). Notably, the cognitive deficits observed in all of these disorders are qualitatively similar to those observed in patients with damage to prefrontal cortex (PFC) (Brown & Marsden, 1990; Cools et al., 2001; Willcutt et al., in press; Barch et al., 2001; Perlstein et al., 2003). Consequently, the overwhelming tendency in the literature is to attribute patient cognitive deficits to dopaminergic dysfunction within PFC. This is a potentially valid attribution, given that selective disruption to prefrontal DA in monkeys gives rise to cognitive deficits that are similar to those observed under full excitotoxic PFC lesions (Sawaguchi & Goldman-Rakic, 1991; Williams & Goldman-Rakic, 1995). However, there is growing evidence that DA dysfunction within the BG alone can be responsible for frontal-like cognitive deficits (Frank et al., 2004; Rinne, Portin, Ruottinen, Nurmi, Bergman, Haaparanta, & Solin, 2000; Collins, Wilkinson, Everitt, Robbins, & Roberts, 2000; Crofts et al., 2001).

Psychopharmacological studies that transiently manipulate the DA system in healthy individuals can potentially provide a more direct way of isolating the cause of DA-related cognitive deficits. In particular, drugs that target the D2 receptor, which is predominantly expressed in BG relative to PFC (Camps, Cortes, Gueye, Probst, & Palacios, 1989), can help determine the specific contributions of the BG. Previous studies with D2 agonists/antagonists in healthy participants have yielded mixed results, with findings of both cognitive enhancement and impairment associated with different doses and task conditions (Kimberg et al., 1997; Peretti, Danion, Kauffmann-Muller, Grangé, Patat, & Rosenzweig, 1997; Muller, von Cramon, & Pollman, 1998; Mehta, Sahakian, McKenna, & Robbins, 1999; Mehta et al., 2001; Kimberg & D'Esposito, 2003; Luciana, Hanson, & Whitley, 2004). For example, Kimberg et al. (1997) found that in tests of executive function, bromocriptine (a D2 agonist) enhanced performance in participants with low working memory span, but impaired performance in those with high span. This provocative result suggests that differences in WM span are partially characterized by differences in baseline DA levels. However, the use of a D2 agonist raises the question of whether these effects are related to DA in the BG rather than PFC (as was originally suggested).

The present research attempts to clarify many of these complex issues, through the use of behavioral studies on a range of cognitive tasks in healthy participants administered two different D2 agents with opposing effects (cabergoline, a D2 agonist, and haloperidol, a D2 antagonist). Cabergoline was chosen over bromocriptine because it has 7 times greater affinity for the D2 receptor while still having low D1 affinity (Ichikawa & Kojima, 2001), and has fewer negative side effects (Stocchi et al., 2003; Corsello et al., 2003; Colao et al., 2000; Biller et al., 1996; Webster, 1994), Similarly, haloperidol has superior in vivo D2 binding when compared to

[1] A version of this chapter has been submitted to the **Journal of Neuroscience** (Frank & O'Reilly, in press)

other agents (Seeman & Kapur, 2000; Kapur et al., 1996), and has 25 times greater affinity for D2 over D1 (Bymaster et al., 1999). These studies were designed to test specific predictions from an existing biologically-based computational model of DA modulation within the BG, and the BG's modulatory role on frontal cortical function (Frank et al., 2001; Frank, in press; O'Reilly & Frank, in press). As described in the next section, the model makes predictions about the D2 drug's effects on learning from positive and negative feedback in procedural learning tasks, and their effects on working memory updating in standard working memory tasks. These predictions were confirmed in our behavioral studies, providing important support for the specific mechanisms proposed in our model, and the more general point that the basal ganglia may play a more important role in cognitive function than is generally appreciated.

5.1 A Mechanistic Account of Basal Ganglia Dopamine Function in Cognition

As described in the previous chapters, recent computational modeling efforts have incorporated various aspects of the BG/DA system toward explicitly simulating their involvement in cognition (Frank et al., 2001; Frank, in press; O'Reilly & Frank, in press). These models suggest that: (a) the BG plays a largely modulatory role in cognition and action, meaning that it does not directly implement any cognitive process but rather modulates function in cortical regions that do so; (b) DA dynamically modulates this already modulatory BG system; and (c) this DA modulation has effects both on learning and the performance of already-learned behaviors. Furthermore, by virtue of interactions with different areas of frontal cortex (Alexander et al., 1986), our model shows how the basal ganglia can participate in a wide range of cognitive functions, from relatively "low-level" tasks such as procedural learning, up to "higher-level" tasks such as working memory and executive function.

Although the principles of this model were described previously, in both the biology and modeling chapters, I summarize the main components of this model that motivated the pharmacological studies. At the most general level, the BG in our model modulates the selection of one of several action plans represented in frontal cortex. (Frank et al., 2001; Frank, in press). This "action-selection" framework leverages existing knowledge of BG involvement in motor control, in which it is thought to selectively facilitate a preferred motor command in premotor cortex while suppressing competing motor programs (Mink, 1996; Jiang, Stein, & McHaffie, 2003). More specifically, two main projection pathways from the striatum go through different BG output structures on the way to thalamus and up to cortex (Figure 5.1). Activity in the **direct** pathway sends a "Go" signal to facilitate the execution of a response considered in cortex, whereas activity in the **indirect** pathway sends a "NoGo" signal to suppress competing responses.

The model also simulates transient changes in DA levels that occur during positive and negative reinforcement, and their differential effects on the two BG pathways. In animals, phasic bursts of DA cell firing are observed during positive reinforcement (Schultz, 2002; Schultz et al., 1997) which are thought to act as "teaching signals" that lead to the learning of rewarding behaviors (Schultz, 2002; Wickens, 1997). Conversely, choices that no not lead to reward are associated with DA dips that drop below baseline (Schultz, 2002; Satoh et al., 2003). Similar DA-dependent processes have been inferred to occur in humans during positive and negative reinforcement (Holroyd & Coles, 2002; Zald et al., 2004). In our models, DA is excitatory on the direct/Go pathway due to a preponderance of D1 receptors on this subset of striatal cells, while it is inhibitory on the indirect/NoGo pathway due to a preponderance of D2 receptors in this subset (Gerfen, 1992; Aubert et al., 2000; Hernandez-Lopez et al., 1997; Hernandez-Lopez et al., 2000). Positive reinforcement leads to DA bursts which increase synaptic plasticity in the direct pathway while decreasing it in the indirect pathway (Nishi et al., 1997; Centonze et al., 2001), supporting "Go" learning to reinforce the good choice. Dips in dopamine have the opposite effect, supporting "NoGo" learning to avoid the bad choice (Frank, in press; O'Reilly & Frank, in press).

This model predicted crossover interaction effects of DA medication on cognitive reinforcement learning in Parkinson patients; as described in the previous chapter, we recently demonstrated strong support for this prediction (Frank et al., 2004). In a probabilistic learning task, all patients and aged-matched controls learned to make choices that were more likely to result in positive rather than negative reinforcement. The difference was

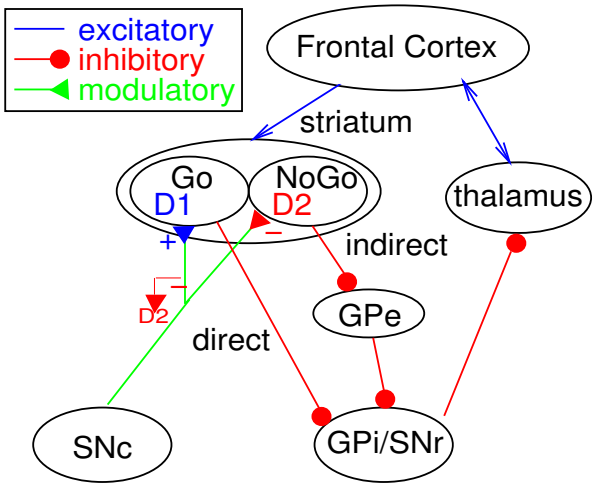

Figure 5.1: The cortico-striato-thalamo-cortical loops, including the direct and indirect pathways of the basal ganglia. The cells of the striatum are divided into two sub-classes based on differences in biochemistry and efferent projections. The "Go" cells project directly to the GPi, and have the effect of disinhibiting the thalamus, thereby facilitating the execution of an action represented in cortex. The "NoGo" cells are part of the indirect pathway to the GPi, and have an opposing effect, suppressing actions from getting executed. Dopamine from the SNc projects to the dorsal striatum, differentially modulating activity in the direct and indirect pathways by activating different receptors: The Go cells express the D1 receptor, and the NoGo cells express the D2 receptor. D2 receptors are also present presynaptically, and provide negative feedback on the amount of DA release during bursting. Dopamine from the VTA projects to ventral striatum and frontal cortex (not shown). GPi: internal segment of globus pallidus; GPe: external segment of globus pallidus; SNc: substantia nigra pars compacta; SNr: substantia nigra pars reticulata.

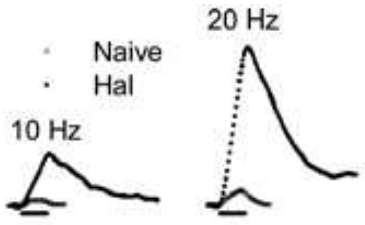

Figure 5.2: Haloperidol administration vastly increases spike dependent DA release upon external DA cell stimulation. This effect is most significant when the cells are stimulated at physiologically relevant frequencies characteristic of bursting. From Wu et al, 2002.

in their strategy: patients taking their regular dose of DA medication implicitly learned more about the positive outcomes of their decisions (i.e., they were better at Go learning than NoGo learning), whereas those who had abstained from taking medication implicitly learned to avoid negative outcomes (better NoGo learning). Aged-matched controls did not differ in their tendency to learn more from the positive/negative outcomes of their decisions. According to the model, the tonically low DA levels in PD patients biased them to learn NoGo to choices that lead to negative outcomes; DA medication reversed this bias, causing better Go learning from positive outcomes.

In addition to these lower-level reinforcement learning effects, our model suggests that these Go/NoGo pathways in the BG can also drive the updating of working memory representations in PFC (O'Reilly & Frank, in press; Frank et al., 2001; Beiser & Houk, 1998). Specifically, Go signals cause PFC to update and maintain current sensory information, while NoGo signals prevent updating, enabling robust ongoing maintenance of previously stored information. This role can potentially explain the effects of D2 receptor agonists (e.g., bromocriptine) on working memory tasks (Kimberg et al., 1997), and other working memory/executive function deficits observed in various patient populations such as Parkinson's disease and ADHD (Brown & Marsden, 1990; Cools et al., 2001; Willcutt, Brodsky, Chhabildas, Shanahan, Yerys, Scott, & Pennington, in press; Barch, Carter, Braver, MacDonald, Noll, & Cohen, 2001; Perlstein, Dixit, Carter, Noll, & Cohen, 2003).

5.1.1 Model Predictions for D2 Drugs: Pre- and Post-synaptic Effects

To make predictions about how D2 drugs will affect performance in our model, we must take into account the two different isoforms of the receptor (D2-S and D2-L), expressed in pre- and postsynaptic neurons, respectively (Uziello et al., 2000). Although cognitive effects of D2 agents are often interpreted in terms of postsynaptic effects (e.g., agonists are thought to mimic the effects of DA whereas antagonists block these effects), evidence from the animal literature suggests that low doses of these agents actually exert their effects primarily via presynaptic mechanisms (Richfield et al., 1989; Schoemaker et al., 1997). Specifically, presynaptic D2 autoreceptors tightly control the level of phasic DA release via (inhibitory) negative feedback (Starke, Gothert, & Kilbinger, 1989; Grace, 1995; Schmitz, Benoit-Marand, Gonon, & Sulzer, 2003). Thus, D2 agonists (e.g., cabergoline) stimulate autoreceptors and diminish the amount of phasic DA release, whereas D2 antagonists (haloperidol) actually increase DA bursting and release in the BG (Moghaddam & Bunney, 1990; Wu et al., 2002; Garris et al., 2003; Figure 5.2). In other words, the effects of D2 agents on phasic DA release is exactly opposite to the pervasive assumptions of these effects in cognitive studies; this is critical for our predictions for drug effects on reinforcement learning processes.

Nevertheless, the drugs should also have direct postsynaptic effects, which should translate into overall behavioral biases on Go/NoGo responding. As noted above, postsynaptic D2 receptors are predominantly expressed in the indirect/NoGo pathway (Gerfen, 1992), which acts to suppress the execution of cortical actions. Because dopamine is inhibitory to D2 receptors in the BG (Hernandez-Lopez et al., 2000), D2 agonists mimic this effect and inhibit NoGo neurons in the indirect pathway (Black, Gado, & Perlmutter, 1997), and should

Drug Mechanism	**Bio Effects** DA	**Behavioral Predictions** Performance	Learning
Cabergoline (agonist)			
D2 post Stim	↑ tonic	↑ Go	
D2 pre Stim	↓ bursts		↓ Go
Haloperidol (antagonist)			
D2 post Block (–)	–	–	
D2 pre Block	↑ bursts		↑ Go

Table 5.1: Summary of Go/NoGo effects of D2 agonists and antagonists, including postsynaptic D2 on overall Go/NoGo activation, and presynaptic D2 effects on DA bursting/release.Due to low dose of haloperidol (2 mg), only presynaptic effects were expected (presynaptic D2 autoreceptors are more sensitive to the drug than postsynaptic receptors).

thereby have a net facilitatory effect on the execution of cortical actions. In other words they should lower the threshold for facilitating responses. In contrast, D2 antagonists should have the opposite effect, disinhibiting the NoGo/indirect pathway and thereby raising the response threshold. This is consistent with the general tendency for DA and its agonists to increase locomotive behavior, while DA blockade leads to catalepsy and Parkinsonism (Fog, 1972).

Thus, taking both pre- and post-synaptic mechanisms into account (Table 5.1), cabergoline should result in more overall Go responding (via postsynaptic stimulation), but less learning of the positive consequences of a given response (via presynaptic reduction of phasic DA release). In contrast, a D2 antagonist such as haloperidol should theoretically enhance NoGo firing (raising the threshold for action and therefore less Go responding) but also enhance DA bursting (thereby enhancing learning of positive outcomes of Go responses). However, due to constraints on safety and risk management, we chose a low dose of haloperidol (2 mg) that is unlikely to have significant postsynaptic effects: while having high affinity for presynaptic autoreceptors, haloperidol only activates postsynaptic receptors at higher doses and/or chronic administration (Schoemaker et al., 1997), at which point catalepsy and Parkinsonism is induced by postsynaptic D2 blockade (Sanberg, 1980).[2] Indeed, unlike patients taking higher doses of the drug (Kumari et al., 1997), participants in our study and other low-dose D2 antagonist studies did not have Parkinson-like slowness of reaction times (Peretti et al., 1997; Mehta et al., 1999). Thus, we predict that the haloperidol effects will hinge on presynaptically-mediated enhancement of DA bursting, supporting increased Go learning.

5.2 Empirical Tests of Unified BG/DA Model

We tested 28 healthy participants (15 females, 13 males) between the ages of eighteen and thirty-five (M 21, SEM 0.75) in a within-subjects design. We conducted three behavioral experiments to test the predictions from our model: probabilistic selection (two alternative forced-choice), probabilistic Go/NoGo learning and reversal, and modified versions of the widely-used AX-CPT working memory task (including set-shifting and reversal conditions). The probabilistic selection task was used in previous work to test our model predictions in PD patients (Frank et al., 2004), and provides a good measure of the effects of dopaminergic changes on implicit learning from positive and negative feedback. The probabilistic Go/NoGo paradigm allows more direct assessment of the Go/NoGo processes that we ascribe to the BG. The AX-CPT working memory task (Servan-Schreiber, Cohen, & Steingard, 1997; Braver, Barch, Keys, Carter, Cohen, Kaye, Janowsky, Taylor, Yesavage, & Mumenthaler, 2001; Barch, Braver, Nystom, Forman, Noll, & Cohen, 1997; Barch et al., 2001) supports the analysis of multiple components of executive function, and is critical for determining whether

[2] Consistent with this, D2 blockade reduced behavioral responding during acquisition of an appetitive response, but actually resulted in **enhanced** conditioned responding in a subsequent test conducted after drug washout (Eyny & Horvitz, 2003) — that is, the drug had a NoGo performance effect during acquisition, but resulted in enhanced Go learning that was only evident once this performance effect wore off.

the same Go/NoGo processes at work in the simpler procedural learning tasks also apply to working memory updating, as predicted by our model.

We also measured participant's working memory (WM) span using the standard reading-span task, as in previous studies (Kimberg et al., 1997), to determine if DA drug effects interacted with baseline WM span. Participants were categorized as low- or high-span according to a median-split on this measure. Each participant was tested on all tasks in three different drug conditions (order randomized): (a) cabergoline (D2 agonist), 1.25mg, (b) haloperidol (D2 antagonist), 2mg, and (c) a placebo. We also measured serum prolactin levels to determine the degree to which the DA drugs had biologically-measurable effects. We begin with the results of these prolactin results, as they constrain our interpretation of the subsequent behavioral results, described thereafter.

5.2.1 Drug Effects on Serum Prolactin Secretion

Serum prolactin levels provide an indirect measure of D2 agent effects, because D2 receptor stimulation inhibits the secretion of this hormone in the pituitary (Ben-Jonathan, 1985). Prolactin levels were obtained before drug ingestion and four hours later, after cognitive tests. There was a main effect of drug on prolactin levels ($F[2,25] = 24.1$, $p < .0001$), such that both active drugs effectively modulated prolactin secretion in opposite directions. Although placebo was associated with diminished prolactin levels over this four hour period (prolactin levels normally decrease during the day), cabergoline significantly decreased prolactin secretion beyond that observed under placebo ($F[1,25] = 44.0$, $p < .0001$), an effect that was observed in both low span ($F[1,25] = 30.0$, $p < .0001$) and high span ($F[1,25] = 14.5$, $p = .0008$) participants. Nevertheless, this effect was somewhat bigger in low span participants, as evidenced by an anti-correlation between prolactin levels under cabergoline and the number of errors made in the working memory span test ($r = -0.38$, $p = 0.08$). In contrast, haloperidol resulted in increased prolactin levels, which was significant relative to placebo ($F[1,25] = 8.9$, $p = .0063$). However, while this effect was significant in low span participants alone ($F[1,25] = 9.0$, $p = .006$), it was not present in those with high span ($F[1,25] = 1.1$, n.s.). Further, the increase in prolactin levels by haloperidol was marginally correlated with the number of errors in the working memory span test ($r = 0.48$, $p = .08$), again suggesting a bigger effect of the drug in low span participants. These span-dependent drug effects on prolactin secretion were unexpected, but provide informative constraints for interpreting similar span-dependent effects on cognition reported below.

5.2.2 Cognitive Task I: Probabilistic Selection, Two Alternative Forced-Choice

In the **Probabilistic Selection** (PS) task (Frank et al., 2004), three different stimulus pairs (AB, CD, EF) are presented in random order and participants have to learn to choose one of the two stimuli (see Methods). Feedback follows the choice to indicate whether it was correct or incorrect, but this feedback is probabilistic. In AB trials, a choice of stimulus A leads to correct (positive) feedback in 80% of AB trials, whereas a B choice leads to incorrect (negative) feedback in these trials (and vice-versa for the remaining 20% of trials). CD and EF pairs are less reliable: stimulus C is correct in 70% of CD trials, while E is correct in 60% of EF trials. Over the course of training participants learn to choose stimuli A, C and E more often than B, D, or F. Note that learning to choose A over B could be accomplished either by learning that A leads to positive feedback, or that B leads to negative feedback (or both). To evaluate whether participants learned more about positive or negative outcomes of their decisions, we subsequently tested them with novel combinations of stimulus pairs involving either an A (AC, AD, AE, AF) or a B (BC, BD, BE, BF); no feedback was provided. If participants had learned more from positive feedback, they should reliably choose stimulus A in all novel test pairs in which it is present. On the other hand, if they learned more from negative feedback, they should more reliably avoid stimulus B.

We predicted that when taking haloperidol, participants should learn more about Go and less about NoGo, due to the enhancement of phasic DA bursts (via presynaptic autoreceptor inhibition) during positive feedback.

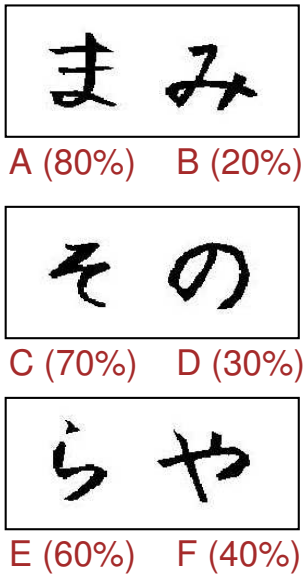

Figure 5.3: Example stimulus pairs for the Probabilistic Selection task. To minimize explicit verbal encoding, the stimuli were Japanese Hiragana characters. Each pair is presented separately in different trials, with participants pressing the "z" key on the keyboard to select the stimulus on the left, or the "m" key to select the stimulus on the right. Three different pairs are presented in random order; correct choices are determined probabilistically. Note that in actuality the position of the correct stimulus was randomized across trials, and the assignment of Hiragana character to hierarchical element was randomized across participants. Different Hiragana characters were used in each session, so these were always novel.

For cabergoline, we predicted an overall Go bias (i.e., speeded reaction times), but impairments in Go learning due to reduced phasic DA release during positive feedback.

5.2.2.1 Results and Discussion

Both drugs displayed the predicted effects on the novel test pair generalizations (Figure 5.4a). We filtered out participants who did not perform above 50% on the most trivial training pair (AB) during test; their test results are meaningless if they could not reliably choose A/avoid B in this pair (Frank et al., 2004). Despite no main effects of drug ($F[2,24] = 2.1$, n.s.) or positive/negative test condition ($F[1,24] = 0.38$, n.s.), the interaction between drug and test condition was highly significant ($F[2,24] = 12.4$, $p = .0002$). Specifically, there was a cross-over interaction between the effects of cabergoline and haloperidol on positive/negative feedback learning ($F[1,24] = 22.7$, $p < .0001$). Compared with placebo, cabergoline ingestion was associated with lower tendency to choose the most "positive" stimulus (i.e., A) in the novel test pairs ($F[1,24] = 4.8$, $p = .038$). In contrast, haloperidol ingestion was associated with more reliable choice of stimulus A than the placebo condition ($F[1,24] = 3.3$, $p = 0.08$). Comparing the two active drug conditions, haloperidol was associated with significantly greater tendency to choose the most positive stimulus ($F[1,24] = 15.6$, $p = .0006$). Nevertheless, haloperidol was associated with **worse** performance in test pairs in which participants had to avoid the most "negative" stimulus (B), compared with both placebo ($F[1,24] = 8.9$, $p = .006$) and cabergoline ($F[1,24] = 11.4$, $p = .002$). These differences in Go/NoGo learning under cabergoline and haloperidol are strikingly similar to the pattern observed in PD patients on and off medication in the same task (Frank et al., 2004) (Figure 5.4b), and support the notion that (a) under cabergoline, decreased DA release during positive feedback impaired Go learning for good choices, and (b) under haloperidol, increased DA release during positive feedback enhanced Go learning for good choices, but may have caused excessive Go learning during positive feedback for stimuli that have overall negative value. Finally, these differences in learning biases on the drugs were found despite no main effect of drug on overall accuracy during training ($F[2,24] = 1.5$, n.s.), and no effect of either cabergoline ($F[1,24] = 2.1$, n.s.) or haloperidol ($F[1,24] = 0.01$), relative to placebo.

For overall Go/NoGo bias effects, we examined reaction times and restricted our analysis to the first 10 trials of the training session, in order to disentangle possible effects due to learning. We reasoned that a Go bias effect should cause speeded RT's, whereas a NoGo bias should cause slowed RT's. As predicted, cabergoline sped up RT's relative to placebo ($F[1,26] = 6.4$, $p = .018$), consistent with an overall Go bias caused by tonic postsynaptic D2 stimulation. Haloperidol was not associated with slowed RT's (as would be predicted by postsynaptic D2 blockade) but was actually associated with numerically faster RT's ($F[1,26] = 1.64$, n.s.). Note that our hypothesis that haloperidol increases DA release during positive feedback predicts that the drug may speed RT's specifically in trials immediately following this feedback, due to elevated DA levels. Indeed, haloperidol was associated with more reliably faster RT's in trials immediately following positive feedback ($F[1,26] = 3.1$, $p = .09$), but not negative feedback ($F[1,26] = 0.25$). In contrast, the tonic Go bias induced by cabergoline was significant no matter whether the previous trial's feedback was positive ($F[1,26] = 6.04$, $p = .02$) or negative ($F[1,26] = 4.6$, $p = .04$).

Span Effects

Although the PS task is not a working memory task, we hypothesized that if baseline working memory span is dictated in part by individual differences in the DA system (Kimberg et al., 1997), then this could potentially be apparent in non-WM tasks, especially given the differences on biological measures we observed with prolactin levels. Interestingly, these span-dependent drug effects were observed even at the level of reaction times. That is, a Go bias effect under cabergoline, in terms of speeded RT's, was observed in low span participants ($F[1,26] = 5.95$, $p = .02$), whereas high span participants did not have significantly faster RT's under cabergoline ($F[1,26] = 1.1$, n.s.). Further, the tendency for haloperidol to speed RT's after positive feedback was near-significant in low span participants alone ($F[1,26] = 3.9$, $p = .06$) but was not evident in those with high span ($F[1,26] = 0.2$). These results are consistent with our observed biological effects of the drugs: while cabergoline decreased prolactin levels in both span groups, this effect was larger in low-span participants, and

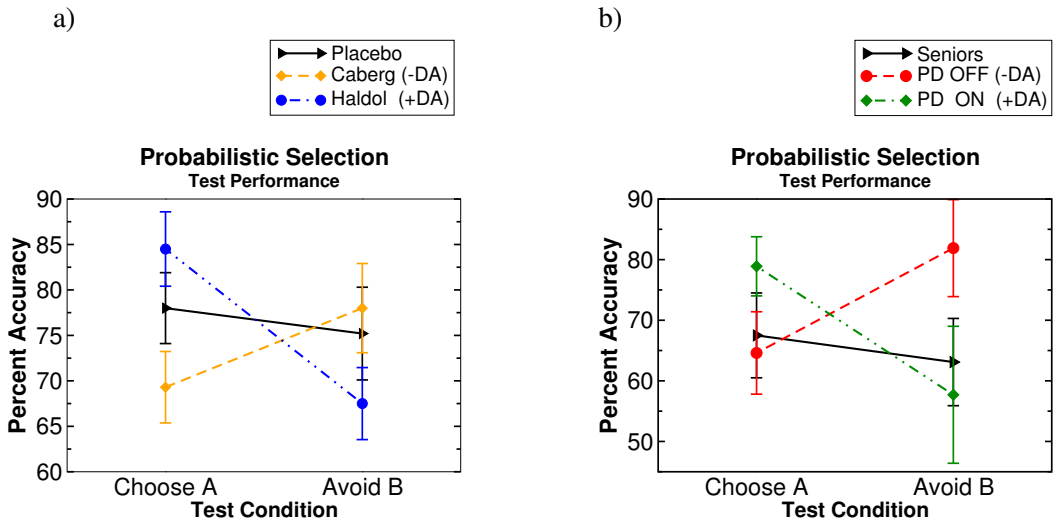

Figure 5.4: Results of probabilistic selection task in (a) healthy participants taking cabergoline and haloperidol (present study); and (b) Parkinson's patients on and off dopaminergic medication (Frank et al, in press). There was a significant interaction between medication condition and tendency to learn from positive vs. negative feedback. Go learning is reflected by increased tendency to choose A in all novel test pairs, whereas NoGo learning is reflected by increased tendency to avoid B. Participants on placebo learned equally to "Go A" and "NoGo B". Haloperidol enhanced positive feedback learning overall, causing a corresponding decrease in NoGo accuracy. Cabergoline impaired positive feedback learning. These results are consistent with our predictions, and mirror the results found previously in Parkinson's patients on and off medication (Frank et al, in press).

haloperidol only significantly increased prolactin levels in low-span participants.

The same general pattern was observed for learning effects. There was no main effect of WM span ($F[1,24] = 0.65$), and no interaction between WM span and positive/negative test condition ($F[1,24] = 0.09$). The effect of WM span on positive/negative learning was only apparent under the drug conditions, in terms of a significant interaction between WM span and drug ($F[2,24] = 3.5$, $p = .045$). Planned contrasts revealed that the tendency for haloperidol to improve positive feedback learning but impair negative feedback learning was significant in low span participants ($F[1,24] = 11.7$, $p = .0022$), but not high span participants ($F[1,24] = 1.15$, n.s.). There was also a significant correlation between positive feedback learning under haloperidol and the number of errors in the working memory span test ($r = 0.51$, $p = .025$), again pointing to a bigger effect of haloperidol in low span participants. No such correlation between positive feedback learning and span errors was observed under placebo ($r = 0.15$, n.s.). For cabergoline, there was no correlation between span errors and positive feedback learning under cabergoline ($r = -.07$, n.s.).

Thus, participants with low WM span were more susceptible to haloperidol enhancements in positive feedback learning than were those with high WM span, while the effects of baseline span on cabergoline effects were less clear. We discuss further significance of these findings later.

5.2.3 Cognitive Task II: Probabilistic Go/NoGo and Reversal

The second task involves a probabilistic reinforcement Go/NoGo paradigm, in which stimuli are presented one at a time and the participant has to either press a key (Go) or withhold their response (NoGo) (see Methods). They are told that some stimulus patterns will give them a point if selected, while others will make them lose a point, and they are to try to maximize point totals. After Go responses, visual feedback is provided ("You won a point!" written in blue or "You lost a point" written in red). Six different patterns are presented in random order, associated with reinforcement probabilities of 80%, 70%, 60%, 40%, 30% and 20% (Figure 5.5. Over time, participants have to learn that three of the stimuli should be associated with a button press (because their corresponding probabilities of reinforcement are greater than 50%), but that responses made to the other three will likely make them lose points. A test session follows the training block, in which the training stimuli and novel combinations of these stimuli are presented without feedback. In novel combinations, the left and right halves of the combined pattern each represent one of the training patterns. For example, half of the composite pattern may consist of a familiar pattern that is 80% correct, while the other half consists of one that is 80% incorrect, so that the combined pattern should be equally associated with "Go" and "NoGo". In some cases one of the patterns is more strongly associated (i.e., 80% combined with 60%), but in others the associations are equal (80/80). Following this test session, a second training block ensues in which the probabilities of reinforcement to the different patterns are reversed. Participants have to learn to respond to stimuli that were previously incorrect and stop responding to stimuli that were previously correct.

We predicted that when taking haloperidol, participants will have learned more about Go responses than NoGo responses (due to enhanced DA bursts during positive feedback), and will therefore tend to respond Go to the combined pattern. Conversely, those taking cabergoline should (a) have a Go bias effect in terms of more overall Go responding, due to tonic postsynaptic D2 stimulation, but (b) should be impaired at learning the positive consequences of these Go responses, due to presynaptic reduction of DA bursts. For reversal effects, we predicted that enhanced DA bursts under haloperidol in the initial training segment would cause excessive Go responding to previously positive stimuli. In contrast, we predicted that cabergoline would lead to deficits in Go learning to previously negative stimuli.

5.2.3.1 Results and Discussion

Training

During training, there was a main effect of drug ($F[2,26] = 3.9$, $p = 0.03$), but no drug by training condition interaction ($F[2,26] = 0.75$, n.s.). Moreover, there were no significant differences when comparing

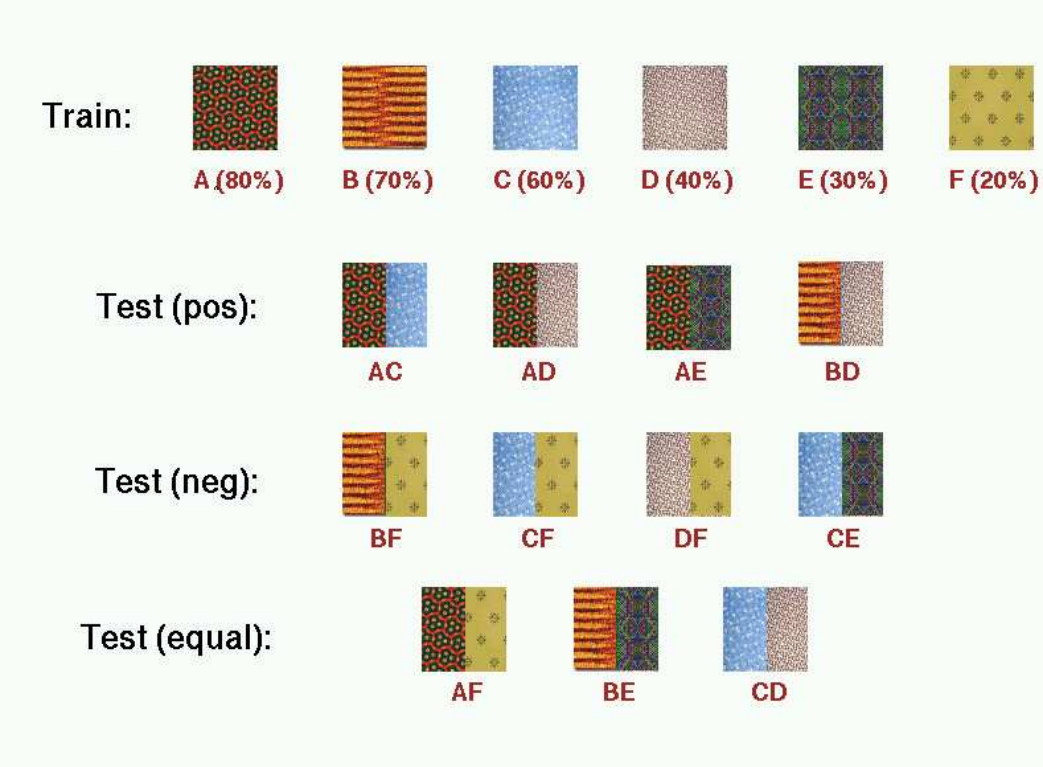

Figure 5.5: .

Example stimuli for the Probabilistic Go/NoGo task. Each training stimulus is presented alone in separate trials. Participants have to either respond “Go” (by pressing the spacebar) or “NoGo” (by withholding their response). After Go responses, feedback is presented with probabilistic reinforcement (percent positive reinforcement indicated in parentheses). Half of the stimuli are associated with greater probability of positive reinforcement, while the other are half are more likely to result in negative reinforcement. During the test phase, training stimuli are re-presented interleaved with novel composites of two training stimuli. In the test(pos) pairs, the combined value of the composite stimulus is positive. In the test(pos) pairs, the combined value of the composite stimulus is negative. In the test(equal) pairs, the combined value of the composite stimulus is neutral (i.e., the individual stimulus elements are equally associated with Go and NoGo responses). Note that in actuality the left and right positions of the stimuli were randomized across trials, and the assignment of texture pattern to frequency of positive feedback was randomized across participants. Different patterns were used in each session, so these were always novel.

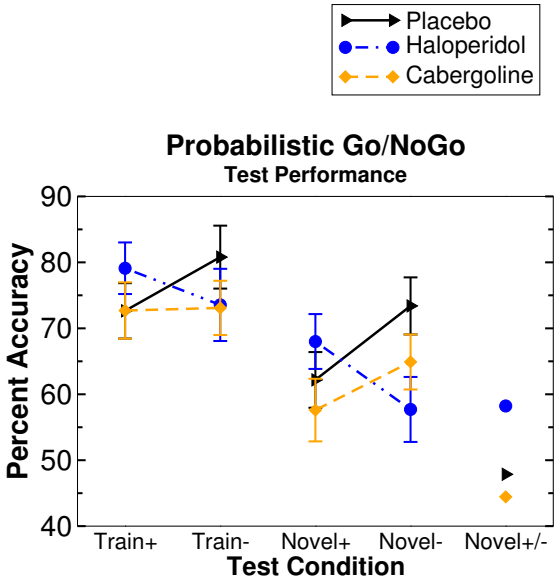

Figure 5.6: Probabilistic Go/Nogo accuracy in a test phase, in which single patterns from the training phase and all possible two-way combinations of these patterns are presented without feedback. Train+: single training stimuli associated with positive reinforcement (pooled 80%, 70% and 60% correct feedback); Train-: training stimuli associated with negative reinforcement (pooled 80%, 70% and 60% incorrect feedback). Because training stimuli can be memorized, a better test of relative Go/NoGo learning comes from the novel test pairs. Novel+: novel combinations of training patterns that together have a more positive than negative association (e.g., the left half of the composite pattern was a 80% Go stimulus, while the right half was a 60% NoGo stimulus); Novel-: novel combinations of training stimuli that have an overall negative association; Novel+/-: combined test pairs that have equal positive and negative associations. Haloperidol resulted in enhanced performance on overall positive pairs and impaired performance on negative pairs, consistent with enhanced Go learning. Cabergoline resulted in a Go bias effect, in terms of more overall Go responding, but was associated with a relative impairment in Go learning for positive stimuli. In balanced Novel+/- trials, haloperidol was again associated with more Go responding as if the novel composite had an overall positive value.

cabergoline ($F[1,26] = 0.1$, n.s.) or haloperidol ($F[1,26] = 1.4$, n.s.) to placebo. Nevertheless, cabergoline was associated with a Go learning deficit: there was a trend for less Go responding across all positive stimuli ($F[1,26] = 2.37$, $p = .14$), which was significant when comparing Go responding for the most positive stimulus (80% correct) ($F[1,26] = 7.66$, $p = .01$). In contrast, haloperidol was associated with numerically but not significantly better positive-Go responding, compared with placebo ($F[1,26] = 1.53$, n.s.). This enhancement in Go responding was not evident in the first block of training ($F[1,26] = 0.96$), but there was a trend for improvement in the second block ($F[1,26] = 2.7$, $p = 0.11$), suggesting that increased Go responding was due to positive feedback learning, rather than an overall propensity for Go responding.

Test The test phase provides a more sensitive measure of implicit Go/NoGo associations (participants could have effectively "memorized" the correct response (Go or NoGo) for each of the six training stimuli). Despite no main effects of drug ($F[2,26] = 1.74$, n.s.) or positive/negative test condition ($F[1,26] = 0.07$), there was a significant interaction between drug and test condition ($F[2,26] = 8.44$, $p = .0015$). For cabergoline, our results support the notion that the drug lowered the overall threshold for making a response (via post-synaptic mechanisms) but impaired the ability to learn positive consequences of these responses (via presynaptic mechanisms). Thus cabergoline increased Go responding overall, as evidenced by a deficit in withholding responses (NoGo) to negatively associated test pairs, relative to placebo ($F[1,26] = 12.1$, $p = .0018$). This deficit can be simply explained by an overall Go bias effect for cabergoline (as was also observed in the PS task in terms of faster RT's). However, it is also possible that the same mechanism responsible for the Go bias also gave rise to a specific impairment in NoGo **learning**: during negative feedback, the drug would continue to bind to postsynaptic D2 receptors and therefore may block the effect of DA dips that would normally facilitate this kind of learning (Frank, in press; Frank et al., 2004). Note that despite this apparent Go bias, cabergoline was not associated with enhanced Go responding for positive test pairs ($F[1,26] = 0.04$, n.s.), suggesting a Go learning deficit. Indeed, compared with haloperidol, cabergoline was associated with significantly less Go responding to positive stimuli ($F[1,26] = 5.8$, $p = .02$).

For haloperidol, the drug increased Go responding, both to positive stimuli ($F[1,26] = 4.97$, $p = .03$) and to those having overall negative value ($F[1,26] = 14.2$, $p = .0008$), compared with placebo. As in the PS task, the haloperidol deficit in withholding responses to negative stimuli can be attributed to "over-counting" the positive feedback that occurs in the minority of trials for these stimuli (due to the probabilistic nature of the task). Because there was no significant effect of haloperidol on Go responding during training (particularly early in training), or on the training pairs that were re-presented at test, the overall enhanced Go responding on novel pairs suggests that participants implicitly learned more from positive than from negative feedback.

For evenly-matched novel pairs (e.g., a 80% Go paired with an 80% NoGo), there was no significant difference between cabergoline and placebo ($F[1,26] = 0.01$), whereas haloperidol was associated with marginally greater Go responding in these cases than placebo ($F[1,26] = 3.5$, $p = 0.07$). Moreover, the difference between the cabergoline and haloperidol conditions was significant ($F[1,26] = 4.5$, $p = 0.04$), with haloperidol causing more Go responding to the composite stimulus as if it had an overall positive value, and cabergoline causing less Go responding as if the composite had an overall negative value.

Reversal In the reversal phase, there was no main effect of drug ($F[2,26] = 0.16$) and no interaction between drug and stimulus condition ($F[2,26] = 0.5$). Nevertheless, a Go learning deficit was apparent under cabergoline, in terms of less Go responding for the most positive stimulus (which was previously the most negative) in the first block of reversal ($F[1,26] = 5.9$, $p = .02$). The predicted haloperidol deficit in withholding responses to previously positive stimuli was evident numerically but was not significant ($F[1,26] = 1.4$, $p = 0.25$). However, we note that this lack of significant effect was not because participants under haloperidol successfully withheld their responses, but rather because all participants (including those under placebo) were unsuccessful at this, as evidenced by overwhelmingly worse tendency to withhold responses to previously positive stimuli, compared to their ability to respond Go to previously negative stimuli ($F[1,26] = 26.1$, $p < .0001$).

Span Effects

As in the PS task, there was no main effect of WM span on test performance ($F[1,26] = 0.6$), and no interaction between WM span and positive/negative test condition ($F[1,26] = 0.3$), consistent with the notion that these tasks do not tax working memory processes. Nevertheless, there was an interaction between the drug effects reported above and baseline WM span ($F[2,26] = 7.1$, $p = .003$), that was very similar to the span interactions observed in the PS task.

First, the overall Go bias effect under cabergoline observed in the PS task in terms of speeded RT's was again evident only in low span participants, who responded Go **more** often than the placebo condition for both positive ($F[1,26] = 5.8$, $p = .024$) and negative stimuli ($F = 10.3$, $p = .003$). Thus this Go bias effect may have masked the Go learning deficits induced by the drug. In stark contrast, those with high WM span (who were not subject to the Go bias in the PS task) responded Go significantly **less** often for positive stimuli under cabergoline ($F[1,26] = 6.4$, $p = .018$).

The same pattern of Go learning deficits were observed in the reversal phase—in terms of less Go responding for the most positive stimulus—in high span ($F[1,26] = 4.6$, $p = .04$), but not low span participants ($F[1,26] = 0.56$), compared with placebo. Again, the Go bias effect under cabergoline may have masked the Go learning deficits in low span participants.

For haloperidol, the drug significantly enhanced Go responding to positive stimuli in low span participants ($F[1,26] = 7.2$, $p = .01$), but not in high span participants ($F[1,26] = 0.15$), mirroring the results of the PS task, and consistent with the drug's effects on prolactin levels and RT's following positive feedback in low span participants.

5.2.4 Summary: Cognitive Procedural Learning

Across both procedural learning tasks, haloperidol resulted in better Go learning from positive feedback than NoGo learning from negative feedback. Conversely, cabergoline resulted in impaired Go learning but more overall Go responding. These results are consistent with the notion that (a) 2mg haloperidol exerts its effects via presynaptic mechanisms and enhances DA bursting; (b) 1.25 mg cabergoline has both pre and postsynaptic effects which mediate the Go learning impairments and baseline Go bias effects, respectively.

Each of the two tasks provides some information not available in the other. The forced choice task demonstrates differential learning from positive and negative feedback which is unconfounded by simple motor effects, since a single motor command is required on each choice trial. However, this task alone does not provide evidence of specific modulation of Go and NoGo learning—theoretically the positive and negative feedback effects could be caused by other (unspecified) neural mechanisms. The converging evidence from a simple Go/NoGo paradigm therefore strongly supports our BG/DA framework.

5.3 Cognitive Task III: Working Memory and Executive Function

In the third task, we test the implications of our model in higher order executive function and working memory (see Methods). To do so, we use a common working memory paradigm called the AX-CPT (continuous performance task) (Servan-Schreiber et al., 1997; Braver et al., 2001; Barch et al., 1997; Barch et al., 2001). We also modified the task to include distractors (Braver et al., 2001), and added reversal and set-shifting conditions. We expected that to the extent the medications modulate responsiveness in the procedural domain, they will similarly affect the updating of working memory in PFC. We used both short (1s) and long (3s) delay conditions, and variable numbers of distractors (0-3) in the long delay (Figure 5.7).

5.3.1 Short Delay AX-CPT

The participant is presented with sequential letter stimuli (A,X,B,Y) (printed in red), and is asked to detect the specific sequence of an A (cue) followed by an X (probe) by pushing the right button. All other cue-probe combinations (A-Y, B-X, B-Y) should be responded to with a left button push. In the short delay

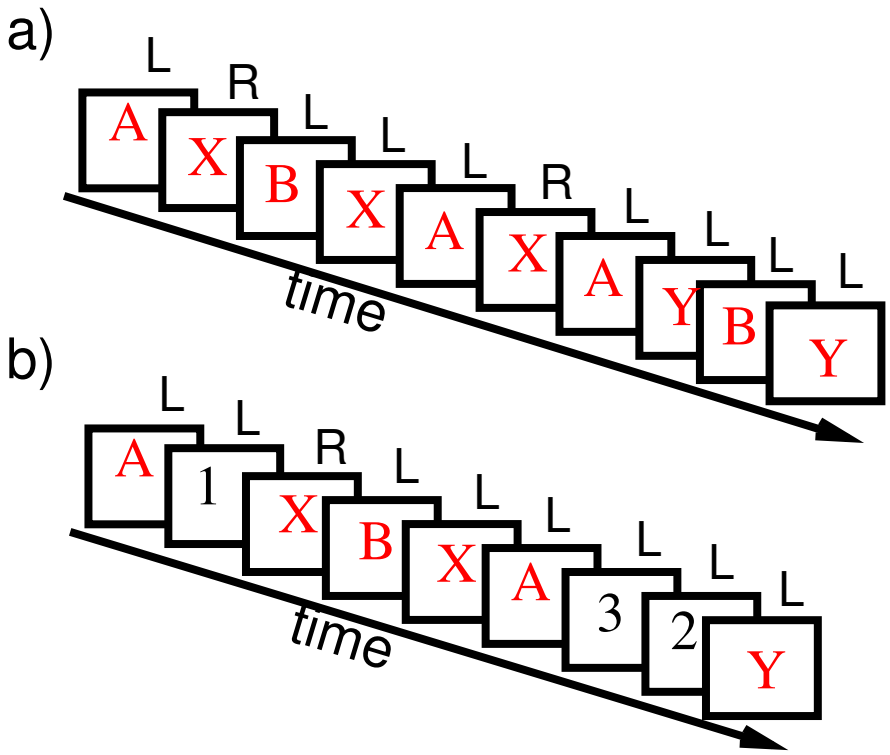

Figure 5.7: The AX-CPT task. (A) Standard version. Stimuli are presented one at a time in a sequence. The participant responds by pressing the right key (R) to the target sequence, otherwise a left key (L) is pressed. Delay between each stimulus is 1 second. The A-X target sequence occurs on 70% of trials, building up a prepotent expectation for target responses. (B) Variable distractors. Task is the same as in (A) but anywhere from zero to three distractors are presented sequentially during a 3 second delay period. Participants have to respond to distractors with a left button push, but are told to ignore these for the purpose of target detection.

case, the delay between cue and probe is one second. The target A-X sequence occurs on 70% of trials, and the other sequences are divided equally by the remaining 30% of trials. This task requires a relatively simple form of working memory, where the prior stimulus must be maintained over a delay until the next stimulus appears, so that one can discriminate the target from non-target sequences. This task also allows analysis of the **type** of errors made. If participants successfully maintain contextual information (e.g., A) in PFC then they will perform well at detecting the A-X target sequence, but will likely make more false positive errors on the A-Y sequence (due to anticipation of an X). Context maintenance also should improve performance on the B-X case, because one can use the B to know not to respond to the X as a target. The B-Y sequence serves as a control, because neither the B or the Y are associated with the target. Furthermore, because the A-X sequence occurs with high probability, it is not as reliable an indicator of working memory performance as B-X and A-Y (good performance on A-X can arise either from learning a prepotent response to stimulus X **or** from maintaining stimulus A in working memory).

Note that participants that do not maintain context can actually do better at A-Y sequences, because they "forget" the A and are not inclined to consider Y as a target. This prediction of improved performance from impaired context maintenance provides a nice counter to a global degradation confound. Nevertheless, these participants should make more errors in the B-X cases (because they have to remember the A/B to decide how to respond to X). This pattern of results was observed in aged participants, presumably due to worse overall PFC context maintenance (Braver et al., 2001).

In our BG/PFC framework, we argue that BG Go signals lead to the updating of PFC working memory representations (Frank et al., 2001), and that DA enhances Go signals, leading to more WM updating. A key prediction of our models is that the level of D2 stimulation modulates the BG threshold for updating working memory in PFC: just as D2 stimulation enhances Go responding in the motor domain, so it should enhance Go signals to update PFC working memory representations (Frank, in press; O'Reilly & Frank, in press). Thus we predicted that cabergoline, by tonically stimulating postsynaptic D2 receptors, would enhance the degree to which WM representations are updated and subsequently maintained. This should translate to better performance on the B-X case, but more A-Y false alarms. For haloperidol, because we had no evidence for postsynaptic D2 blockade, we did not predict that WM representations would be less updated, but rather that enhanced DA bursts for positive stimuli (in a WM task, task-relevant stimuli are "positive" (O'Reilly & Frank, in press)) would also lead to more WM updating. Importantly, the different mechanisms for these two drug effects are dissociated by predictions in the modified versions of the task, described later.

5.3.1.1 Results and Discussion

To examine working memory effects, we computed a **working memory context index** by subtracting BX - AY performance; increased maintenance should improve BX performance but impair AY performance, and vice-versa for decreased maintenance (Figure 5.8). A positive WM context index indicates greater influence of working memory on choice behavior, whereas a negative context index indicates that choices are being dictated by incoming stimuli and are not influenced by working memory (Braver et al., 2001). That this index is a reliable measure of working memory is supported by a significant anti-correlation with number of errors in the working memory span test ($r = -0.48$, $p = .01$).[3] . There was a main effect of drug on WM context ($F[2,26] = 6.2$, $p = .0063$), such that cabergoline significantly enhanced WM ($F[1,26] = 12.3$, $p = .0016$), and haloperidol marginally enhanced WM ($F[1,26] = 3.8$, $p = .06$), relative to placebo. Moreover, the context index was significantly positive under both cabergoline ($t[1,26] = 4.71$, $p < .0001$) and haloperidol ($t[1,26] = 2.72$, $p = .01$), whereas it did not differ significantly from zero under placebo ($t[1,26] = 0.36$, n.s.)

To ensure that drug effects were genuinely related to working memory, we analyzed performance in the B-Y control trials, in which performance does not depend on working memory. Importantly, there was no main effect of drug on performance in these trials ($F[2,26] = 0.94$, n.s.), and no effect of either cabergoline ($F[1,26]$

[3] The placebo condition was used for the context index in this analysis, since the working memory span test was only performed in pre-screening sessions without drugs

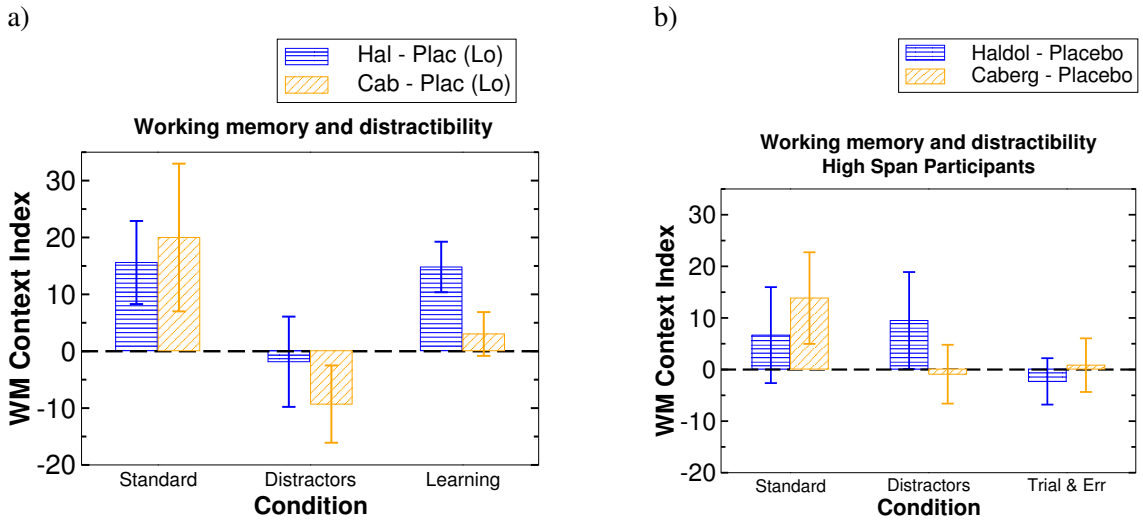

Figure 5.8: Within-subject (relative to placebo) AX working memory results as a function of drug and task condition, broken down into low and high span participants. The WM context index = BX - AY indicates the degree to which working memory influences performance (see text for explanation). **(a)** Low span participants. In the standard condition (short delay), both drugs were associated with increased WM context relative to placebo. Under placebo, working memory context generally decreased with longer delay and distractors (not shown). Cabergoline resulted in further increased distractibility, whereas haloperidol did not. These results are consistent with our theory that cabergoline lowered the overall threshold for updating working memory (and therefore also caused distractibility), whereas haloperidol selectively increased updating for task-relevant information via enhanced DA bursts. In the procedural version of the A-X task, feedback allows participants to learn target sequences by trial-and-error by reinforcing responses and associated working memory representations. In this case, haloperidol increased WM context, again consistent with an enhancement of DA release during positive feedback. Counteracting learning and Go bias effects for cabergoline explain a lack of enhancement/impairment under this drug in the procedural A-X task. **(b)** Drug effects were generally smaller or nonexistent for high span participants. For haloperidol, this is consistent with a lack of effect on prolactin levels, speeded RT's and positive/negative reinforcement learning in these participants. For cabergoline, this is consistent with smaller effects on prolactin levels and a lack of Go bias effect in procedural learning tasks.

= 0.23) or haloperidol (F[1,26] = 1.86, n.s.), relative to placebo.There was also no main effect of drug on A-X target responding (F[2,26] = 1.3, n.s.), and no effect of either cabergoline (F[1,26] = 0.65) or haloperidol (F[1,26] = 2.6, n.s.), relative to placebo.

Span Effects

There was no interaction between baseline WM span and drug effect (F[2,26] = 0.18) in the short-delay AX-CPT.

5.3.2 Long Delay AX-CPT With and Without Distractors

The short-delay case supported our general hypothesis that tonic postsynaptic D2 stimulation by cabergoline enhances BG Go signals and therefore leads to more updating of PFC WM representations. Note, however, that this same hypothesis predicts that cabergoline should also cause concomitant updating to task-irrelevant information, and therefore cause increased distractibility. This reasoning allows us to make distinct predictions for tonic (cabergoline) vs. phasic (haloperidol) effects on BG DA signals. Specifically, tonic D2 stimulation by cabergoline should result in counteracting effects of more updating of useful information, but also more updating of distractors. In contrast, a purely phasic increase in DA as produced by haloperidol should cause enhanced updating selectively for task-relevant information (but see below for possible negative consequences of this in set shifting), assuming that our model is correct in claiming that DA bursts occur only for stimuli that should be maintained for a given task (and are therefore "positive").

To test this prediction, we introduce distractor stimuli (single white digits [1,2,3,4]) during the cue-probe delay interval. The number of distractors between the cue and target varies randomly from zero to three. We also increased the delay between cue and probe to three seconds, to allow enough time to present up to three distractors.

5.3.2.1 Results and Discussion

For working memory effects, there was a large main effect of distractors (F[1,26] = 10.4, p = .003), such that the WM context index was significantly decreased by distractors, as expected. There was no main effect of drug (F[2,26] = 0.5, n.s.), but there was a significant drug by distractor interaction (F[2,26] = 4.1, p = .029). Notably, the effect of cabergoline interacted significantly with distractor condition (F[1,26] = 8.0, p = .009). Relative to placebo, cabergoline resulted in increased WM context index in zero distractor trials (F[1,26] = 3.3, p = .08) but decreased WM when distractors were present (F[1,26] = 4.5, p = .04). Moreover, in zero-distractor trials the WM context index was significantly positive under cabergoline (F[1,26] = 5.06, p = .033), while it did not differ reliably from zero under placebo (F[1,26] = 0.09). In the distractor conditions, however, cabergoline was associated with a significantly negative WM context index (F[1,26] = 7.67, p = 0.01), indicating that working memory for task-relevant information was impaired, while it still did not differ reliably from zero under placebo (F[1,26] = 0.15). These results support our claims that cabergoline enhanced working memory updating and concomitant distractibility. Note that the decreased context index under distractors does not translate to global performance decrements. In fact, whereas B-X performance decreased with distractors, A-Y performance actually improved under cabergoline—the updating of distracting information may have led to forgetting of the A and thus less anticipatory A-Y false alarms. Further, just as in the short delay case, there was no effect of cabergoline on B-Y control conditions (F[1,26] = 0.24) or A-X target sequence performance (F[1,26] = 0.25).

For haloperidol, there was no interaction between drug and distractors (F[1,26] = 1.06, n.s.). In zero-distractor trials, the WM context index was numerically, but not significantly greater than placebo (F[1,26] = 1.5, n.s.), and was marginally greater than zero (F[1,26] = 2.8, p = 0.10). When distractors were present, WM context was not different from that under placebo (F[1,26] = 0.07) or from zero (F[1,26] = 0.5). Further, there were no differences between haloperidol and placebo in either in B-Y (F[1,26] = 0.03) or A-X performance (F[1,26] = 0.63). Thus, haloperidol improved WM context in the short delay case and was associated with a

marginally positive context index in long delay zero-distractor trials, without causing increased distractibility. This is consistent with the idea that haloperidol increased phasic Go bursting for task-relevant information, causing it to be better stored in PFC, but was not associated with increased Go signals for task-irrelevant information (which are unlikely to be associated with DA bursts). We test more specific predictions for haloperidol to selectively improving WM under conditions related to phasic DA signals in our novel "Procedural A-X Task" described later.

Span Effects There was no overall interaction between baseline WM span and overall drug effect ($F[2,26] = 1.3$, n.s.). Nevertheless, planned contrasts revealed that the differential effect of cabergoline to enhance/impair WM context depending on distractors was significant for low span participants ($F[1,26] = 9.4$, $p = .005$), but was not present in those with high span ($F[1,26] = 0.5$, n.s.). This is because while cabergoline significantly enhanced WM context for the zero distractor condition in low span participants ($F[1,26] = 7.3$, $p = .01$) it did not in those with high span ($F[1,26] = 0.16$). As in the procedural learning tasks, the Go bias effect induced by cabergoline may have been present only in low span participants, a claim that is also supported by the stronger drug effect on prolactin levels in these participants.

5.3.3 Reversal and Set-Shifting

As in the Go/NoGo procedural learning task, we were also interested in the effects of the drugs on reversal and set-shifting. These processes are typically associated with prefontal function, but they also require changing the criteria for updating task-relevant versus irrelevant information. It is exactly this kind of dynamic changing of updating thresholds that we think the BG/DA system is specialized to perform (Frank, in press, under revisions; O'Reilly, Noelle, Braver, & Cohen, 2002; Rougier & O'Reilly, 2002). For the reversal case, we reversed the target sequence from A-X to B-Y. For the set-shifting case, we swapped the targets and distractors, making the target sequence 1-3 (vs. 2-3, 1-4, 2-4), while the letters A, X, B and Y became distractors. A final reversal switched the target to 2-4.

We predicted that on top of the primary effects on updating and distractibility listed above, cabergoline should specifically impair the simple reversal so that participants would continue to respond to A-X sequences as if they were targets during the B-Y target block. The prepotent tendency to respond to these A-X sequences, given their high frequency in both the previous short and long delay segments, should be particularly difficult to override with an additional cabergoline Go bias effect. Further, this perseverative responding should only occur in zero-distractor A-X trials: when distractors are present, cabergoline should continue to increase WM updating and should therefore lead to **less** anticipation of A-X sequences and less perseveration. No effect of haloperidol was expected in reversal, because our hypothesized effect for the drug to enhance DA bursting to update task-relevant information does not change in this case (i.e., the [A,X,B,Y] letters that were task-relevant in the A-X segment continue to be task-relevant in the B-Y segment).

In the set-shifting segments, we predicted impairments caused by both drugs, but for different reasons. For cabergoline, a Go learning deficit should translate into difficulty in learning to update newly task-relevant information. Thus overall deficits should be observed when participants have to learn to now update the white digit stimuli [1,2,3,4] which used to be distracting. In contrast, we predicted that enhanced phasic DA signals under haloperidol should not result in difficulty in updating new information, but should impair set-shifting when previously task-relevant stimuli become distractors. That is, during the initial A-X and B-Y segments the enhanced DA bursts will have established a strong Go updating signal to these task-relevant letter stimuli, which will cause increased distractibility in conditions where these letters are now distractors.

5.3.3.1 Results and Discussion

In the reversal segments (i.e., when the targets switched from A-X to B-Y, or from 1-3 to 2-4), there was no main effect of drug on the ability to respond to new target sequences ($F[2,26] = 0.7$), and no drug interaction with distractors ($F[2,26] = 0.35$). There was also no main effect of cabergoline ($F[1,26] = 0.07$)

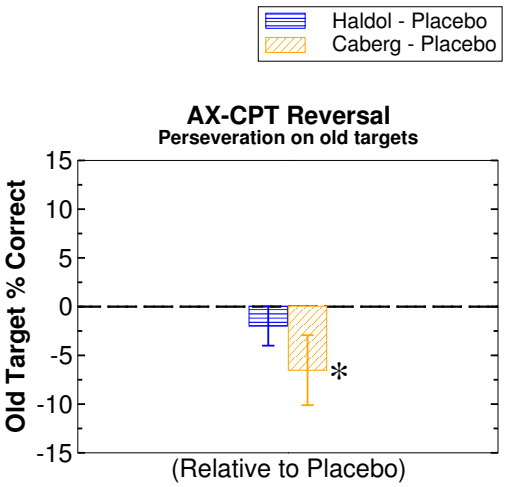

Figure 5.9: Reversal of target sequences from A-X to B-Y and 1-3 to 2-4 resulted in perseverative responding to old target sequences (A-X and 1-3) under cabergoline (relative to placebo). This impairment is consistent with postsynaptic D2 stimulation by the drug, causing a lowered threshold for executing prepotent responses. The plot shows performance on old target sequences in zero-distractor trials, for which previous-target anticipation is strongest. Span effects are not plotted due to the small number of participants who actually perseverated (most participants performed at ceiling for previous targets).

or haloperidol ($F[1,26] = 0.8$), when compared to the placebo condition. There were deficits, however, in perseverative responding to old target sequences under cabergoline. More specifically, these deficits interacted with distractors ($F[1,26] = 6.8$, $p = .01$), such that perseverative responding to previous target sequences was observed selectively in zero-distractor trials , with **less** perseveration than placebo under distractors. That this is a reversal deficit is particularly compelling, given that the A-X sequence is effectively the "control" condition for the B-Y segment; when analyzed across all segments there was no main effect of cabergoline on control condition performance ($F[1,26] = 0.24$) and no interaction with distractors ($F[1,26] = 0.09$). Further, there were no effects of haloperidol on control conditions, either across all segments (main: $F[1,26] = 0.03$; distractor interaction: $F[1,26] = 1.0$, n.s.) or in the reversal segments (main: $F[1,26] = 0.02$; distractor interaction: $F[1,26] = 0.05$).

The reversal deficits under cabergoline are consistent with the hypothesis that the Go bias effect causes participants to anticipate a target response upon seeing stimulus A (since the A-X sequence was prepotent in the previous segment). But when distractors are present, cabergoline may cause updating of these distractors and therefore less target anticipation.

Reversal Span Effects

There was no main effect of span on perseveration in the reversal segments ($F[1,26] = 0.09$), and no interaction between drug and span ($F[2,26] = 1.36$, n.s.).Nevertheless, under cabergoline, there was an interaction between perseveration and distractor condition for low span participants ($F[1,26] = 8.2$, $p = .008$), who perseverated on the old target sequence in the reversal segment significantly more than placebo ($F[1,26] = 4.7$, $p = .039$), while there was a trend for **better** performance than placebo in distractor trials ($F[1,26] = 3.0$, $p = .095$). There was no such interaction between perseveration and distractors for high span participants ($F[1,26] = 0.39$, n.s.), who did not differ from placebo either in the zero distractor condition ($F[1,26] = 0.0$.) or under distractors ($F[1,26] = 1.05$, n.s.), again consistent with a lack of a cabergoline Go bias effect in these participants.

In the set-shifting (1-3) segment, where participants must now ignore the letters and maintain numbers, there were impairments for both cabergoline and haloperidol (Figure 5.11). For cabergoline, there were global performance decrements in the set-shifting segment ($F[1,26] = 5.0$, $p = .033$), but no interaction with distractors ($F[1,26] = 0.24$). This suggests that shifting deficits under cabergoline were due to an inability to update the new set, rather than selective deficits that may relate to ignoring the old set. Indeed, deficits under cabergoline were marginally significant in just zero-distractor trials ($F[1,26] = 3.35$, $p = .079$), with non-significant deficits under distractors ($F[1,26] = 1.9$, n.s.). Thus cabergoline deficits in set-shifting are consistent with our hypothesis that reduced DA bursts for newly task-relevant ("positive") stimuli resulted in Go learning deficits to update the new set.

In contrast, overall set-shifting deficits under haloperidol ($F[1,26] = 5.24$, $p = .03$) interacted with distractor condition ($F[1,26] = 3.9$, $p = .058$). In particular, there was no difference between haloperidol and placebo in zero-distractor trials ($F[1,26] = 0.06$), but there was a large performance decrement when the previously task-relevant letter stimuli became distractors ($F[1,26] = 10.5$, $p = .0032$). These results are consistent with our hypothesis that enhanced DA bursts supported working memory updating of the task-relevant letters in the initial stages. But while this may have been useful for differentiating relevant versus distracting information and therefore resulted in reduced distractibility, it is detrimental to performance when it comes time to ignore previously important information. That these are set-shifting specific deficits is suggested by analyzing just the two segments before the set-shift: in this analysis there was no main effect of haloperidol ($F[1,26] = 0.24$), and no interaction with distractors ($F[1,26] = 0.56$) on global performance measures.

Set-Shifting Span Effects

In the set-shifting segment, there was no main effect of span ($F[1,26] = 2.0$, n.s.), but the interaction between drug and span was significant ($F[2,26] = 4.15$, $p = .027$).Specifically, under cabergoline, low span participants were not impaired in either zero-distractor ($F[1,26] = 0.13$) or distractor ($F[1,26] = 1.1$, n.s.) trials. In contrast, high span participants were selectively impaired in the zero distractor condition ($F[1,26] = 9.9$, $p = .004$), with no significant impairments for distractors ($F[1,26] = 0.78$). In other words, high span participants

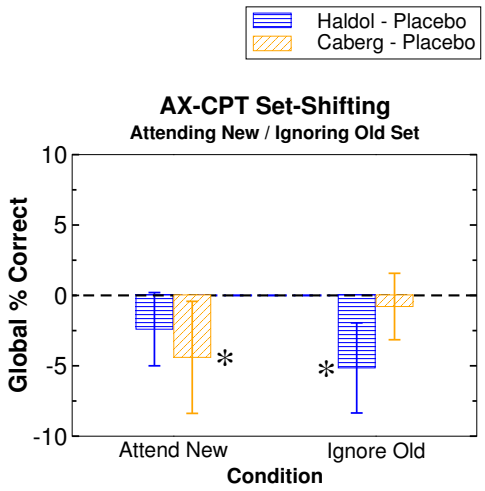

Figure 5.10: Set-shifting results in the AX-CPT working memory task. Drug effects are shown within-subject relative to the placebo condition, in terms of global performance across all trial types. Cabergoline resulted in significant impairments in attending to the new but previously distracting set, as evidence by performance in the zero distractor condition of the 1-3 set-shifting segment. Haloperidol was not associated with significant set-shifting deficits in the zero distractor condition, but resulted in difficulty ignoring previously task-relevant letters when these letters subsequently became distractors. This is consistent with the hypothesis that the drug enhanced DA bursts for the letter stimuli in these participants during the initial A-X and B-Y segments, making them difficult to subsequently ignore.

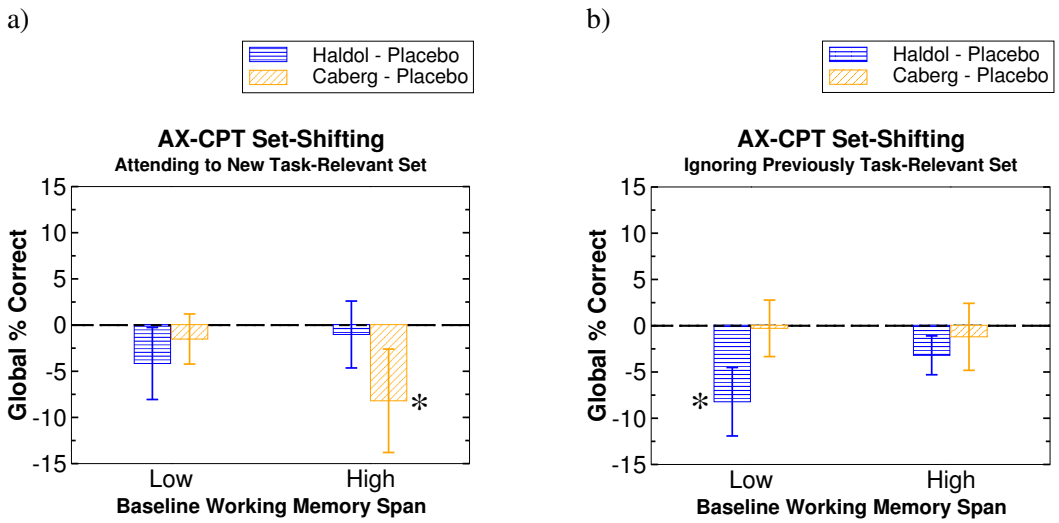

Figure 5.11: Set-shifting results in the AX-CPT working memory task, broken down into low and high span participants and drug effects relative to placebo. **a)** Set-shifting deficits under cabergoline, in terms of difficulty attending to a previously distracting set (white digits), were observed in high span participants. These deficits were evident by global performance decrements in the zero-distractor condition during the initial (1-3) set-shifting segment. That they were observed only in high span participants is consistent with the Go learning deficit induced by cabergoline in these participants in both procedural learning tasks. In low span participants, the Go bias effect to update all stimuli may have masked an otherwise Go learning deficit. Haloperidol was not associated with set-shifting deficits in the zero distractor condition. **b)** Set-shifting deficits under haloperidol, in terms of difficulty ignoring previously task-relevant letters, when these letters subsequently became distractors. This is consistent with the hypothesis that the drug enhanced DA bursts for the letter stimuli in these participants during the initial A-X and B-Y segments, making them difficult to subsequently ignore. As in procedural learning tasks, haloperidol effects were only observed in low span participants.

were impaired at attending to the new set, because their impairments were observed when only stimuli from this set were presented. For haloperidol, the set-shifting interaction with distractor condition was marginally significant in low span participants ($F[1,26] = 3.9$, $p = .06$), who had worse performance in distractor trials ($F[1,26] = 12.15$, $p = .0018$), with no difference between drug conditions in zero-distractor trials ($F[1,26] = 0.23$). No such interaction was observed in high span participants ($F[1,26] = 0.57$, n.s.) participants, who did not differ from placebo either in zero-distractor trials ($F[1,26] = 0.03$) or under distractors ($F[1,26] = 0.88$).

Thus, low span participants under haloperidol had deficits in ignoring previously task-relevant distractors, while high span participants under cabergoline had deficits in attending to the new set. These results parallel those observed in the procedural learning tasks, in which low span participants displayed enhanced Go learning under haloperidol, presumably due to bigger DA bursts. These DA bursts should make it possible to attend to a new set, but may explain why these participants were impaired at ignoring the previously relevant set. In contrast, high span participants under cabergoline exhibited impaired Go learning in the Go/NoGo task, and impaired Go responding to previously negative stimuli in the reversal phase, which may explain their impairments in attending to a new set in the set-shifting segment. That is, diminished DA bursts may have resulted in impoverished BG Go signals needed to update new task-relevant information into prefrontal WM representations. For low span participants, the cabergoline Go bias effect observed in both procedural learning and WM tasks may have counteracted this Go learning deficit and enabled them to update the new set.

5.3.4 Procedural A-X

In the above working memory tasks, the participant is instructed which target sequence (e.g., AX) to respond to before each segment. To provide more direct comparison with the procedural learning tasks, we included a novel version of A-X in which participants have to figure out the target sequence by trial-and-error. As in the above AX tasks, letters (H,K,Z,P) are presented sequentially one at a time. No distractors are present (as this segment is more difficult in and of itself). Participants are instructed to press the left button for each cue, and the right button when they think they have seen the target sequence (initially by guessing). After each probe stimulus, feedback informs the participant if they were correct or incorrect. This "procedural A-X" task may depend more on phasic DA signals modified by the drugs than the standard working memory tasks. Further, this version may correspond better to WM tasks used in non-human primate experiments, in which animals have to learn which working memory representations to reinforce via reward or lack thereof.

We predicted that due to enhanced DA bursts during positive feedback, haloperidol should increase WM context in this task. In contrast, cabergoline should result in counteracting Go bias effects but diminished DA bursts during positive feedback, making it difficult to predict the effect of this drug (no distractors were included in this task).

5.3.4.1 Results and Discussion

In the trial-and-error version of the AX-CPT, there was a main effect of drug on the WM context index ($F[2,26] = 3.6$, $p = .04$). Planned contrasts revealed that while cabergoline did not enhance WM context relative to placebo ($F[1,26] = 0.2$), haloperidol did so significantly ($F[1,26] = 6.7$, $p = .016$). This finding further supports our hypothesis that haloperidol enhanced DA bursts, which should be particularly apparent in WM tasks that depend on trial-and-error. In contrast, cabergoline had no effect on WM context, which our framework reconciles by our claim that the drug has two effects, a presynaptic inhibition of DA bursts, which should impair WM context in WM tasks that depend on these bursts (such as trial and error or reversal), and a postsynaptic Go bias effect that lowers the overall threshold for updating working memory, which should enhance WM context, as in the standard version of the task.

As in the standard version of the tasks, there were no main effects of either cabergoline ($F[1,26] = 0.01$) or haloperidol ($F[1,26] = 1.0$) on B-Y control conditions. For A-X target sequences, there was no effect of cabergoline ($F[1,26] = 0.65$), but there was a trend for worse performance under haloperidol ($F[1,26] = 2.9$, p

= 0.1).

Span Effects

There was no main effect of span on WM context in the procedural A-X task ($F[1,26] = 0.3$), but there was a trend for an overall interaction between span and drug condition ($F[2,26] = 2.4$, $p = 0.11$). Planned contrasts revealed that the enhancement of WM context under haloperidol was significant in low span ($F[1,26] = 9.98$, $p = 0.004$) but not high span participants ($F[1,26] = 0.15$). These results are again consistent with the effect of haloperidol to increase DA bursts in low span but not high span participants. Under cabergoline, neither those with low span ($F[1,26] = 0.2$) or high span ($F[1,26] = 0.02$) differed in WM context from placebo.

5.4 General Discussion

Taken together, the results of our studies support an emergent, unified framework for basal ganglia dopamine function that cuts across divergent cognitive domains. This framework goes beyond descriptive theories emphasizing **that** the BG system is involved in a particular cognitive task, to attempt to explicitly address **how** DA modulation within the BG affects cognition. By adopting this mechanistic approach, motivated through explicit computational simulations, our framework reconciles findings of cognitive enhancements and impairments that interact with drug, dose, and task conditions. Specifically, our framework accounts for findings across all three experimental tasks (and several sub-conditions within each task) in terms of the following effects of the D2 agonist cabergoline and the D2 antagonist haloperidol.

In brief, our framework suggests that dopamine dynamically modulates BG "Go" and "NoGo" signals which facilitate and suppress the execution of cortical actions. To make predictions for the effects of D2 agents, we consider both pre- and postsynaptic effects of D2 receptor stimulation and blockade in the BG. Acute administration of low doses of these drugs act primarily on highly sensitive presynaptic D2 autoreceptors (Schoemaker et al., 1997), which tightly control the level of DA release via inhibitory feedback. Therefore, despite pervasive assumptions for D2 agonists/antagonists to stimulate/block the dopamine system in general, these drugs actually have the opposite effects on phasic DA release. By stimulating D2 autoreceptors, cabergoline actually diminishes the amount of DA release; haloperidol enhances DA release during bursting (Wu et al., 2002; Garris et al., 2003, Figure 5.2). In our models, these effects translate to cabergoline impairing, and haloperidol enhancing, the learning of choices that are likely to lead to positive reinforcement. Nevertheless, by acting on postsynaptic D2 receptors, the drugs should also modulate the overall threshold for executing responses. Thus tonic postsynaptic D2 stimulation by cabergoline should produce an overall bias on Go responding similar to that produced by enhanced overall dopamine levels, despite Go learning deficits resulting from decreased DA release during positive reinforcement.

A summary of the primary results found in this study provides overwhelming support for the general predictions of these models (Table 5.2). For cabergoline, the drug greatly reduced prolactin levels and likely exerted its cognitive effects via both pre- and postsynaptic D2 stimulation. Thus cabergoline was associated with an overall Go bias (speeded RT's and more overall Go responding) via presumed postsynaptic stimulation, but impaired Go learning from positive feedback via presynaptic inhibition of DA bursts. Interestingly, these same effects extended to working memory (WM) and attentional domains. In working memory, the overall Go bias may have also enhanced Go signals that support the internal updating of working memory representations. Thus, the drug improved WM specifically in task conditions for which it is always good to update WM (i.e., when all stimuli were task-relevant). But when distracting information was presented during the delay period, the lowered WM updating threshold may have increased distractibility, such that cabergoline actually impaired WM maintenance of task-relevant information. Finally, the Go learning deficits observed under cabergoline in procedural learning tasks translated into impairments in attending to a new set in set-shifting conditions (i.e., in learning Go to update stimuli that were previously distracting).

The effects of haloperidol were consistent with a primarily presynaptic blockade of highly sensitive D2 autoreceptors (Schoemaker et al., 1997), causing enhanced dopamine release during bursting and leading to

relative to placebo	**Cabergoline** post D2 Stim	pre D2 Stim	**Haloperidol** post D2 Block	pre D2 Block
Biological Drug Effects				
Prolactin	$\downarrow\downarrow$		$\uparrow$	
Effective DA	$\uparrow$ tonic^{ls}	$\downarrow$ bursts	–	$\uparrow$ bursts^{ls}
Procedural Learning				
Performance	$\uparrow$ Go^{ls}		–	
Learning		$\downarrow$ Go		$\uparrow$ Go^{ls}
Reversal	$\uparrow$ Go^{ls†}	$\downarrow$ Go$_{New}{}^{hs}$		$\uparrow$ Go$_{Old}{}^{ls*}$
Executive Function				
Working memory Updating	$\uparrow$ Go^{ls}		–	$\uparrow$ Go$_{task-relevant}{}^{ls}$
Distractibility	$\uparrow$ Go^{ls}		–	
Reversal	$\uparrow$ Go$_{prevtarg}{}^{ls}$		–	
Set-Shifting		$\downarrow$ Go$_{New}{}^{hs}$		$\uparrow$ Go$_{Old}{}^{ls}$
Learning (procedural A-X)	$\uparrow$ Go†	$\downarrow$ Go†	–	$\uparrow$ Go^{ls}

Table 5.2: Summary of behavioral results in procedural learning (performance, learning, reversal) and executive function (working memory updating, distractibility, reversal and set-shifting), and how these may relate to drug mechanisms of tonic and phasic D2 stimulation/blockade. Null effects that correspond to null biological effects are indicated by "–". Hypothesized biological drug effects on D2 pre- and postsynaptic receptors are constrained by their observed effects on prolactin (PRL), and by existing knowledge on the greater presynaptic sensitivity. ^{ls} Effect only observed in participants with low working memory span, who had greater biological drug effects, in terms of prolactin secretion. ^{hs} Effect only observed in high span participants, due to selective Go learning deficits without concomitant overall Go bias effects. * This observed effect was not statistically significant, potentially related to the tendency even for participants under placebo to continue to respond Go to previously-positive stimuli. † Counteracting pre- and postsynaptic effects may have led to a null behavioral effect.

increased Go learning from positive feedback. Postsynaptic D2 receptor antagonism effects such as slowed RT's and less Go responding (as are generally observed under higher doses of haloperidol; Kumari et al., 1997), were not observed with our 2 mg dose. Rather, RT's were actually speeded following positive feedback, and Go learning was enhanced in procedural learning tasks. Again, these same effects extended to working memory and attentional domains: haloperidol selectively enhanced working memory updating of task-relevant (i.e., "positively-valenced"), but not task-irrelevant ("negatively-valenced") information. These results support suggestions that DA bursts (which are enhanced by haloperidol) occur not only in response to overt positive reinforcement, but also for task-relevant working memory updates (O'Reilly & Frank, in press; O'Reilly et al., 2002; Rougier & O'Reilly, 2002; Braver & Cohen, 2001). Further, haloperidol also enhanced working memory context in our novel procedural A-X task — in which participants had to learn working memory target sequences via positive/negative reinforcement — presumably due to enhanced DA bursts during positive feedback. However, we do not suggest that low doses of haloperidol should always improve executive function. In fact, the drug was associated with performance decrements during set-shifting conditions when previously task-relevant stimuli were presented as distractors. Thus, too much DA release during the initial segments may have caused excessive attention (i.e., Go) to these stimuli when they should have been subsequently ignored.

On top of the primary effects reported above, we also found some effects to depend on baseline working memory span, which were not explicitly predicted by our model. Overall, both cabergoline and haloperidol had larger effects in participants with low WM-span, in terms of both behavioral effects and the biological measure of prolactin secretion. For cabergoline, the overall Go bias in both procedural learning and executive function tasks was only evident in low span participants. Nevertheless, Go learning deficits under cabergoline were still observed in high span participants, raising the possibility that the drug had a selective presynaptic stimulation effect (i.e., diminished DA bursting) in these participants. For haloperidol, both biological (prolactin) and behavioral effects in the present study were only observed in low-span participants.

Why should D2 drugs have larger effects in low-span participants? Most accounts of span-dependent drug effects leverage the idea that proficient WM is associated with an optimal level of DA in PFC (Williams & Goldman-Rakic, 1995; Arnsten, 1997). This account suggests that DA agonists enhance performance in low span participants by increasing their low DA levels, but may cause excessive DA stimulation in those with high span who already have an optimal DA level (Kimberg et al., 1997; Mattay, Callicott, Bertolino, Heaton, Frank, Coppola, Berman, Goldberg, & Weinberger, 2000). The present results point to a somewhat different account. First, we suggest that the basal ganglia is a more relevant locus of D2 drug effects rather than PFC (see below). Second, we suggest that biological differences between low and high span participants may include differential sensitivity to D2 receptor stimulation, rather than overall differences in dopamine levels. However, it may be possible to reconcile these two accounts, by suggesting that low tonic levels of dopamine in low-span participants leave the D2 receptor more susceptible to influence from the D2 drugs, whereas the higher levels of dopamine in high-span participants cause the drug effects to be relatively diluted through receptor binding competition.

5.4.1 Relationship to Other Studies

Our findings are consistent with those of several other existing studies using dopaminergic agents. First, the opposite effects of cabergoline and haloperidol on Go/NoGo learning directly replicate our prior results in Parkinson's patients, which were reversed by DA medication (Frank et al., 2004, Figure 5.4). Diminished DA bursts in the BG during positive feedback, caused by either PD or by cabergoline administration, resulted in better learning from negative than positive feedback. Increased release of BG DA, caused either by L-Dopa medication in PD or by haloperidol administration, resulted in better learning from positive feedback. Second, our cabergoline results showing reversal learning impairments are consistent with similar impairments in participants taking bromocriptine (also a D2 agonist) (Mehta et al., 2001). Notably, in that study bromocriptine also improved performance in a spatial working memory task. Because that task only involved task-relevant stimuli, this improvement is also consistent with cabergoline effects in our study.

As noted earlier, span-dependent bromocriptine effects on working memory have also been reported (Kimberg et al., 1997). In that study, low span participants improved in set-shifting, while high span participants were impaired. Our findings suggest that these results may have stemmed from differential receptor effects of the agonist across span groups. Specifically, the overall Go bias effect in low span participants may have resulted in enhanced working memory updating overall. In contrast, selective reduction of DA bursts in high span participants may have caused deficits in updating new task-relevant information in set-shifting conditions. Consistent with this account, in the present study we observed Go learning impairments in reversal and set shifting conditions only in the high-WM-span participants on cabergoline. More generally, our findings that dopaminergic agents have span-dependent effects even on "lower-level" cognitive tasks is consistent with a recent animal study showing that low working memory performance was predictive of greater locomotor response to amphetamine (Dellu-Hagedorn, in press).

Our haloperidol results are consistent with animal studies showing that single low doses of D2 antagonists enhance DA release (Moghaddam & Bunney, 1990; Wu et al., 2002; Garris et al., 2003). In both animals and humans, the phenomenon of latent inhibition is disrupted with enhanced DA release (e.g., Gray, Pickering, Hemsley, Dawling, & Gray, 1992). This effect was also recently observed with low dose D2 antagonists (Barrett, Bell, Watson, & King, 2004), and is consistent with enhanced Go learning observed in the present study. Further, the set-shifting deficits under haloperidol in our study are consistent with similar set-shifting impairments under sulpiride, another D2 antagonist (Mehta et al., 1999). Interestingly, although Mehta et al. (1999) interpreted their results as evidence for sulpiride simulating the cognitive profile of PD patients, we emphasize that blocking D2 receptors in healthy participants may actually increase DA release during bursts. This contrasts with the lack of available DA associated with Parkinson's disease. Therefore, we suggest that the observed set-shifting impairments were instead due to presynaptic enhancement of DA bursts before the shift. This is consistent with the fact that no slowed RT's were observed in that study, which would be consistent with a postsynaptic effect. Nevertheless, higher dose antagonists should have significant postsynaptic effects, resulting in an overly active NoGo pathway, and may come closer to simulating PD.

5.4.2 Are D2 Agents Selective for the Basal Ganglia?

Although we have restricted our analysis of D2 drug effects in terms of their preferential action in the BG, we cannot completely discount possible direct drug effects in PFC and other areas. However, we believe this simplification is valid, for several reasons. Anatomically, D2 receptors are by far most prevalent in the BG (e.g., 11 times greater in striatum than in frontal cortex) (Camps et al., 1989). In animal studies, enhanced DA release by haloperidol has been found to be selective in the BG, with DA levels in frontal cortex relatively unaffected (Moghaddam & Bunney, 1990). Further, D2 agents affect working memory processes when applied systemically, but **not** when directly infused into PFC (Yang & Seamans, 1996; Arnsten et al., 1994; Arnsten et al., 1995), suggesting that they exert their effects in the BG. In contrast, D1 receptor stimulation/blockade directly in PFC have profound effects on executive processes (Durstewitz & Seamans, 2002; Sawaguchi & Goldman-Rakic, 1991; Williams & Goldman-Rakic, 1995). In a recent study, Wang, Vijayraghavan, and Goldman-Rakic (2004) showed that blocking PFC D1 receptors reduced working memory related activity, whereas blocking D2 receptors only reduced movement related activity occurring at the end of a delay period, with no reported effect on behavioral performance. Further, we note that if these D2-dependent effects in PFC played a substantial role in the present study then we would have expected haloperidol to cause slowness of responding (by suppressing movement related frontal activity); this was not observed.

5.4.3 Relation to Other Theories of Basal Ganglia Function

The main aspects of our framework are well accepted by several researchers. First, the notion that the BG facilitates/suppresses the execution of cortical actions via separate Go/NoGo pathways is by no means novel (DeLong, 1990; Mink, 1996; Berns & Sejnowski, 1998; Frank et al., 2001). Our suggestion that dopamine

modulates the relative balance of activity in these two pathways is also shared by others; indeed it is often used to account for various motor deficits in patients with neurological DA dysfunction (Albin, Young, & Penney, 1989). The main contribution of our computational models has been to integrate these effects toward exploring their interaction with changing cognitive task demands. Thus, these models simulate phasic changes in DA that are thought to occur during reinforcement learning (Holroyd & Coles, 2002; Schultz, 2002; Satoh et al., 2003), and their differential effect in modifying synaptic plasticity (Nishi et al., 1997; Centonze et al., 2001) in the Go and NoGo pathways. The net effect is that DA bursts during positive feedback reinforce Go representations, whereas DA dips during negative feedback reinforce NoGo representations (Frank, in press; Brown, Bullock, & Grossberg, 2004b). These dynamic processes allow the BG to subsequently facilitate cortical actions that are on average good choices, while concurrently suppressing those that are maladaptive. Finally, our implication of BG/DA in higher level cognitive function is consistent with many who emphasize a critical role of circuits linking the BG with prefrontal cortex (Alexander et al., 1986; Middleton & Strick, 2000b, 2002). In our models, BG Go signals reinforce the updating of PFC working memory (WM) representations; DA enhances these Go signals and effectively lowers the threshold for updating WM (Frank et al., 2001; Frank, in press; O'Reilly & Frank, in press).

5.4.4 Conclusions

In summary, our results support a unified account of the role of DA in modulating cognitive processes that depend on the basal ganglia. Although this account is undoubtedly simplistic (i.e, it does not consider critical interactive effects between DA and other neurotransmitters), it has clear implications that may allow more in-depth understanding of the neural bases for cognitive disorders in Parkinson's disease, ADHD, and schizophrenia. We offer the following cautious interpretation. In Parkinson's disease, low levels of both tonic and phasic DA are associated with diminished updating of PFC working memory representations and less overall frontal activity. In ADHD, low levels of tonic DA may actually result in enhanced DA release during bursting, due to less presynaptic feedback control (Grace, 2001; Seeman & Madras, 2002). According to our framework, this hypersensitivity to phasic DA bursts would cause impulsive behavior and distractibility, due to a lack of control over BG Go signals. Finally, in schizophrenia elevated levels of BG DA, combined with low levels of PFC DA (Seeman, 1987; Weinberger, 1987) may result in too low of a threshold for updating working memory representations, causing attention to inappropriate thoughts, but less PFC maintenance of these representations. Future research with behavioral paradigms such as those used in the present study may help to test these ideas.

5.5 Methods

Procedures were approved by the Scientific Advisory Committee of the University of Colorado Health Sciences Center, and by the University of Colorado Human Research Committee. We used a within-subjects double blind design. Participants reported to the Boulder GCRC for lab tests and a medical exam. A GCRC physician conducted the medical exam, drug test, and reviewed the medical data for each participant. Those who met the study criteria and who received medical approval then proceeded to the pre-experimental session, including the working memory reading span test (Daneman & Carpenter, 1980). They were then scheduled for all three experimental sessions, separated by two weeks.

On the day of the experimental sessions, smoking and drinking alcohol or caffeine was prohibited (enforced via a breathalyzer test). Baseline pre-drug blood samples (5 mL) were drawn and sampled for serum prolactin. Subsequently, participants received a tablet of either 2 mg haloperidol, 1.25 mg cabergoline, or placebo. The assignment of drug type to experimental session number was counterbalanced across participants, as was the order of tasks within a given session. Participants then waited for 2.5 hours to allow the drug to be absorbed (peak plasma levels are reached between 2 and 3 hours for both cabergoline and haloperidol, (Persiani, Rocchetti, Pacciarini, Holt, Toon, & Strolin-Benedetti, 1996; Darby, Pasta, Dabiri, & Mosbacher, 1995)). Dur-

ing this time, participants watched a video. Pulse rates were also monitored every 15 minutes. The cognitive tests lasted approximately 1.25 hours. Four hours after drug ingestion (approximately 10-15 minutes following test completion) a second blood sample was drawn to measure drug effects on serum prolactin levels. Ratings of subjective arousal, mood, restlessness and other side effects were taken pre and post each experimental session using a visual analog scale (Appendix C).

5.5.1 Sample

Our sample was 28 healthy participants, 15 females and 13 males, between the ages of eighteen and thirty-five. Mean age was 21. Exclusion criteria included the use of any medications (prescription or non-prescription), illicit drugs, more than 5 cigarettes a day or 4 cups of coffee per week.

5.5.2 Baseline Working Memory Span

In the pre-screening session, all participants performed the working memory span test (Daneman & Carpenter, 1980). We scored the total number of errors in this test, and then divided the participants in two groups by median split. The median number of errors was 25.

5.5.3 Probabilistic Selection Task Procedures

Participants sit in front of a computer screen in a lighted room and view pairs of visual stimuli that are not easily verbalized (Japanese Hiragana characters, see Figure 5.3). These stimuli are presented in black on a white background, in 72 pt font. They press keys on the left or right side of the keyboard depending on which stimulus they choose to be "correct". Note that the forced-choice nature of the task controls for any differences in overall motor responding. Visual feedback is provided (duration 1.5 seconds) following each choice (the word "Correct!" printed in blue or "Incorrect" printed in red). If no response is made within four seconds, the words "no response detected" are printed in red.

Three different stimulus pairs (AB, CD, EF) are presented in random order. Feedback follows the choice to indicate whether it was correct or incorrect, but this feedback is probabilistic. Choosing stimulus A leads to correct (positive) feedback in 80% of AB trials, whereas choosing stimulus B leads to incorrect (negative) feedback in these trials. CD and EF pairs are less reliable: stimulus C is correct in 70% of CD trials, while E is correct in 60% of EF trials. Over the course of training participants learn to choose stimuli A, C and E more often than B, D, or F.

We enforced a performance criterion (evaluated after each training block of 60 trials) to ensure that all participants were at the same performance level before advancing to test. Because of the different probabilistic structure of each stimulus pair, we used a different criterion for each (65% A in AB, 60% C in CD, 50% E in EF)[4] . After reaching this criterion, participants were subsequently tested with the same training pairs, in addition to all novel combinations of stimuli, in random sequence. They were instructed (prior to the test phase) to use "gut instinct" if they did not know how to respond to these novel pairs. Each test pair was presented 6 times for a maximum of four seconds duration, and no feedback was provided.

5.5.4 Probabilistic Go/NoGo and Reversal Task Procedures

The stimuli employed for this task are colored and textured patterns that are difficult to verbalize (Figure 5.5). As in the selection paradigm, individual patterns that appear in any given session are not reused in other sessions. Each trial begins with a green fixation circle in the center of the screen for 1 second, followed

[4] In the EF pair, stimulus E is correct 60% of the time, but this is particularly difficult to learn. We therefore used a 50% criterion for this pair simply to ensure that if participants happened to "like" stimulus F at the outset, they nevertheless had to learn that this bias was not going to consistently work.

by a stimulus pattern lasting 1 second. For each stimulus, participants press a key on the keyboard or withhold their response. Visual feedback (1 second duration) is provided only after Go responses ("You won a point!" written in blue or "You lost a point" written in red, together with their running "batting average"). If they do not respond within 1 second, feedback written in black letters reads "No points won or lost". Six different patterns (A-F) are presented in random order, associated with reinforcement probabilities of 80%, 70%, 60%, 40%, 30% and 20%. Over time, participants have to learn that three of the stimuli should be associated with a button press (because their corresponding probabilities of reinforcement are greater than 50%), but that responses made to the other three will likely be incorrect.

As in the PS task, we enforced a performance criterion (evaluated after each training block of 60 trials). Participants had to perform at least 70% on stimuli A and F, 60% on B and E and F, and not worse than chance on C and D (which were only 60% reliable). The participant advanced to the test session if all these criteria were met, or after four blocks (240 trials) of training.

In the test session, there are 69 trials in which the training patterns A through F are re-presented 6 times each, interleaved with novel combinations of elemental patterns (AC, AD, AE, AF, BD, BE, BF, CD, CE, CF, DF; that is, four overall Go patterns and four overall NoGo patterns) presented 3 times each. In these novel patterns, the left and right halves of the combined pattern each represent one of the training patterns. For example, half of the composite pattern may consist of a familiar pattern that is 80% correct, while the other half consists of one that is 80% incorrect, so that the combined pattern should be equally associated with "Go" and "NoGo". In some cases one of the patterns is more strongly associated (i.e., 80% combined with 60%), but in others the associations are equal (80/80). Each novel test pattern is presented three times, interleaved among elemental training stimuli that are presented six times each.

5.5.4.1 Reversal

A second training block of 240 trials follows the Go/NoGo test session, in which the probabilities of reinforcement to the different patterns are reversed. Participants have to learn to respond to stimuli that were previously incorrect and stop responding to stimuli that were previously correct.

5.5.5 Working Memory Task Procedures

The working memory segment begins with the standard task with no distractors and an inter-stimulus interval of one second. Each stimulus is presented for 500 ms. There are 50 sequences in the standard task, with 70% A-X target sequences, and 10% each of A-Y, B-X and B-Y.

Next, a longer (3 second) delay period is used during which we present from 0 to 3 distractors sequentially. Each stimulus is presented for 500 ms, and each distractor is presented for 333 ms. The distractors are spaced out evenly throughout the 3 second delay period. When one distractor is presented, there is a 1333 ms delay between the cue and the distractor and between the distractor and the target. For two distractors, the delay between each item is 778 ms. For three distractors, it is 500 ms. Participants have to respond to each distractor with a left button push to ensure that they encode them, but are told to ignore them for the purpose of target detection (Braver et al., 2001). In the A-X distractor segment, there are 32 trials of which 20 are A-X target sequences, and the remaining 12 are divided equally between A-Y, B-X, and B-Y. This same procedure is followed for the reversal and set-shift conditions, for a total of 128 trials in the distractor conditions.

Finally, in the procedural A-X task, different letter stimuli are used (e.g., H,G,Z,P) and participants are instructed to try to figure out what the target sequence is by trial-and-error. They have to respond with a left button push to all cues, and a right button push to the cue-probe combination that they think is a target. Each stimulus is presented for 500 ms, and the delay between cue and probe is 2 seconds. Feedback follows each probe, either "Correct!" written in blue or "Incorrect" written in red, as in the PS task. There are 24 trials, but in this case the target sequence is equally likely to appear as other sequences (a majority of target sequences would make the learning of targets trivial). This is repeated three times with new target sequences, for a total

of 72 trials in the procedural A-X segment.

5.5.6 Analysis

Prior to statistical analysis, we filtered out participants who did not satisfy global performance measures during the test sessions. In the Probabilistic Selection task, we filtered out data from participants who did not perform better than chance (50%) at the most trivial training pair (AB) during test in a given session (Frank et al., 2004). Thus this amounted to six participants on placebo, six on haloperidol and two on cabergoline, who became globally confused by the test phase; their results in choosing among novel pairs are meaningless. Similarly, in the Go/NoGo task we filtered out participants who did not respond above chance on either of the two most trivial training stimuli (80% Go or 80% NoGo). Using these criteria, two participants on cabergoline and one participant on haloperidol were excluded. For the A-X working memory sessions, we filtered out participants who performed worse than 50% on A-X target sequences, which occur on the majority of trials and are trivial – this amounted to one participant on placebo and three on haloperidol, who clearly were not paying attention to the task rules.

In order to be consistent across all data analyses, we performed the same statistical test for each analysis. We used SAS v8.0 PROC MIXED to examine both between and within subject differences. In all analyses we controlled for session number and baseline span group effects, while also testing for interactions between drug and span. In the procedural learning tasks, another factor of positive/negative test condition was added in, along with interactions between this factor and drug and span. In the working memory segments, a distractor factor and its interaction terms (distractor*span, distractor*drug, distractor*span*drug) were included. Where indicated, we tested for specific planned contrasts.

Chapter 6

Conclusions

This dissertation presented a theoretical framework for understanding the interactions between the basal ganglia and frontal cortex in cognitive processes, and how these are dynamically modulated by dopamine. The theory proposes that the BG modulates the strength of representations in cortex by either facilitating (Go) or suppressing (NoGo) cortical commands. Thus the fundamental function of the BG is to gate the execution of commands that are being "considered" in different parts of frontal cortex. Increased levels of dopamine results in more Go and less NoGo, whereas decreases have the opposite effect. The functional implications are summarized for three of the five BG-FC circuits first described by Alexander et al. (1986):

- Motor circuit: the BG modulates learning and execution of stimulus-response associations in premotor cortex, and dopamine results in overall more Go responding. Responses that result in reinforcement elicit phasic increases in DA which transiently activate the direct/Go pathway and suppress the indirect/NoGo pathway. Responses that do not result in reinforcement elicit phasic DA dips which have the opposite effect, transiently activating NoGo representations. Hebbian learning in the two BG pathways supports implicit and procedural learning so that responses that are often reinforced are subsequently facilitated (Go), whereas those that are not reinforcing are suppressed (NoGo). Therefore, participants with lower levels of dopamine release (either due to Parkinson's disease or via acute administration of cabergoline) are more likely to learn about negative consequences of their decisions, whereas those with higher levels of dopamine (induced by L-Dopa medication in Parkinson's patients and by an acute low dose of haloperidol in healthy participants) learn more about the positive consequences. In addition to the presynaptic modulation of DA release, the effects of cabergoline were also consistent with tonic postsynaptic stimulation, leading to a "Go bias" – faster RT's and more overall Go responses (despite impaired learning from their positive consequences).

- Prefrontal circuit: the dorsal caudate gates the learning and maintenance of working memory items to be stored in dorsolateral PFC. Phasic changes in BG DA result in dynamic gating that enables the flexible switching of attentional sets. When the current focus of attention is task-relevant, DA bursts in the BG promote the updating of relevant information via BG Go signals. DA dips after updating irrelevant information may induce NoGo learning so that this information is subsequently ignored. In our pharmacological studies, enhanced DA release by haloperidol resulted in greater updating of task-relevant information, but difficulty ignoring this information when it subsequently became distracting in a set-shift (i.e., there was excessive Go reinforcement in the first segment of the task). In contrast, diminished DA release by cabergoline resulted in difficulty learning to update a new task-relevant set which had previously been distracting (diminished Go learning). Further, tonic postsynaptic D2 stimulation by cabergoline led to overall more working memory updating (Go bias), which improved performance when all stimuli were task-relevant, but caused excessive distractibility when distractors were introduced.

- "Limbic" circuit: the ventral striatum gates the learning and maintenance of stimulus-reward information into orbitofrontal cortex. Because the OFC then maintains the magnitude of expected future

reward, it can then exert top-down control to override responses that would lead to smaller short-term rewards and is therefore implicated in motivation and impulsiveness. The mechanism for this overriding may involve OFC inhibition of DA bursts in the ventral striatum that would otherwise result in facilitation of prepotent responses. This hypothesis is currently being investigated computationally.

In all of these circuits, both tonic and phasic levels of dopamine modulate the threshold for gating. Low levels of DA in Parkinson's disease and chronic drug abuse results in motor hesitancy, low motivation, and lack of working memory updating. In ADHD, oversensitivity to phasic bursts of DA results in hyperactive behavior, impulsiveness, and excessive working memory updating which increases distractibility. Across all conditions, a healthy dynamic range of DA is critical in order to learn changing reinforcement contingencies associated with stimuli, responses, and attention over time.

A considerable amount of cognitive behavioral, neuropsychological and pharmacological observations can be accounted for by this framework. Damage to dopaminergic systems within the BG and PFC produce cognitive deficits that are consistent with this model. However, there is no doubt that the proposed framework ignores several important factors. First, it largely ignores the contribution of a host of other critical neurotransmitters, including acetylcholine, norepinephrine, serotonin and their interactions with dopamine. Second, in the interest of space, I have left out the undoubtedly important role of the anterior cingulate, which in addition to the motor, prefrontal and orbitofrontal circuits, is also under modulation of the BG. Third, while this dissertation adopts the popular hypothesis that DA in the PFC supports robust maintenance of information over time, the focus was more on the role of DA in the BG. This was not meant to underscore the role of DA in PFC, but merely to explore the complementary aspects of BG DA which had not been given as much attention. The hope is that we can better understand the roles of each region by taking a systems-level neuroscience approach.

6.1 References

Abi-Dargham, A., Rodenhiser, J., Printz, D., Zea-Ponce, Y., Gil, R., Kegeles, L. S., Weiss, R., Cooper, T. B., Mann, J. J., Van Heerturn, R. L., Gorman, J. M., & Laruelle, M. (2000). Increased baseline occupancy of D2 receptors by dopamine in schizophrenia. Proceedings of the National Academy of Sciences, 14, 8104–9.

Agid, Y., Ruberg, M., Hirsch, E., Raisman-Vozari, R., Vyas, S., Faucheux, B., Michel, P., Kastner, A., Blanchard, V., Damier, P., Villares, J., & Zhang, P. (1993). Are dopaminergic neurons selectively vulnerable to Parkinson's disease? Advances in Neurology, 60, 148–164.

Aizman, O., Brismar, H., Uhlen, P., Zettergren, E., Levet, A. I., Forssberg, H., Greengrad, P., & Aperia, A. (2000). Anatomical and physiological evidence for D_1 and D_2 dopamine receptor colocalization in neostriatal neurons. Nature Neuroscience, 3, 226–230.

Akkal, D., Burbaud, P., Audin, J., & Bioulac, B. (1996). Responses of substantia nigra pars reticulata neurons to intrastriatal D1 and D2 dopaminergic agonist injections in the rat. Neuroscience Letters, 213, 66–70.

Albin, R. L., Young, A. B., & Penney, J. B. (1989). The functional anatomy of basal ganglia disorders. Trends in Neurosciences, 12, 366–375.

Alexander, G., Crutcher, M., & DeLong, M. (1990). Basal ganglia-thalamocortical circuits: Parallel substrates for motor, oculomotor, "prefrontal" and "limbic" functions. In H. Uylings, C. Van Eden, J. De Bruin, M. Corner, & M. Feenstra (Eds.), The prefrontal cortex: Its structure, function, and pathology, Vol. 85 (pp. 119–146). Amsterdam: Elsevier.

Alexander, G. E. (1987). Selective neuronal discharge in monkey putamen reflects intended direction of planned limb movements. Experimental Brain Research, 67, 623–634.

Alexander, G. E., & Crutcher, M. D. (1990a). Functional architecture of basal ganglia cicuits: neural substrates of parallel processing. Trends in Neuroscience, 13, 266–271.

Alexander, G. E., & Crutcher, M. D. (1990b). Preparation for movement: Neural representations of intended direction in three motor areas of the monkey. Journal of Neurophysiology, 64, 133–50.

Alexander, G. E., DeLong, M. R., & Strick, P. L. (1986). Parallel organization of functionally segregated circuits linking basal ganglia and cortex. Annual Review of Neuroscience, 9, 357–381.

Amalric, M., Moukhles, H., Nieoullon, A., & Daszuta, A. (1995). Complex deficits on reaction-time performance following bilateral intrastriatal 6-OHDA infusion in the rat. European Journal of Neuroscience, 7, 972–980.

Amos, A. (2000). A computational model of information processing in the frontal cortex and basal ganglia. Journal of Cognitive Neuroscience, 12, 505–519.

Armony, J. L., Servan-Schreiber, D., Cohen, J. D., & LeDoux, J. E. (1997). Computational modeling of emotion: explorations through the anatomy and physiology of fear conditioning. Trends in Cognitive Sciences, 1, 28–34.

Arnsten, A. F. (1997). Catecholamine regulation of the prefrontal cortex. Journal of Psychopharmacology, 11, 151–62.

Arnsten, A. F., Cai, J. X., Murphy, B. L., & Goldman-Rakic, P. S. (1994). Dopamine D_1 receptor mechanisms in the cognitive performance of young adult aged monkeys. Psychopharmacology, 116, 143–151.

Arnsten, A. F., Cai, J. X., Steere, J. C., & Goldman-Rakic, P. S. (1995). Dopamine D_2 receptor mechanisms contribute to age-related cognitive decline: the effects of quinpirole on memory and motor performance in monkeys. Journal of Neuroscience, 15, 3429–3439.

Ashby, F. G., Alfonso-Reese, L., Turken, A. U., & Waldron, E. M. (1998). A neuropsychological theory of multiple systems in category learning. Psychological Review, 105, 442–481.

Ashby, F. G., Maddox, W. T., & Bohil, C. J. (2002). Observational versus feedback training in rule-based and information-integration category learning. Memory & Cognition, 30, 666–77.

Ashby, F. G., Noble, S., Ell, S. W., Filoteo, J. V., & Waldron, E. M. (2003). Category learning deficits in Parkinson's Disease. Neuropsychology, 17, 133–42.

Ashby, F. G., Queller, S., & Berretty, P. T. (1999). On the dominance of unidimensional rules in unsupervised categorization. Perception and Psychophysics, 61, 1178–1199.

Atallah, H. E., Frank, M. J., & O'Reilly, R. C. (2004). Hippocampus cortex and basal ganglia: insights from computational models of complementary learning systems. Neurobiology of Learning and Memory, 82(3), 253–267.

Aubert, I., Ghorayeb, I., Normand, E., & Bloch, B. (2000). Phenotypical characterization of the neurons expressing the D1 and D2 dopamine receptors in the monkey striatum. Journal of Comparative Neurology, 418, 22–32.

Baddeley, A. D. (1986). Working memory. New York: Oxford University Press.

Bannon, M. J., Grace, A. A., Bunney, B. S., & Roth, R. H. (1980). Evidence for an irreversible interaction of bromocryptine with central dopamine receptors. Naunyn Schmiedebergs Arch Pharmacol, 312, 37–41.

Bao, S., Chan, V. T., & Merzenich, M. M. (2001). Cortical remodelling induced by activity of ventral tegmental dopamine neurons. Nature, 412, 79–82.

Bao, S., Chan, V. T., Zhang, L. I., & Merzenich, M. M. (2003). Suppression of cortical representation through backward conditioning. Proceedings of the National Academy of Sciences, 100, 1405–1408.

Barch, D. M., Braver, T. S., Nystom, L. E., Forman, S. D., Noll, D. C., & Cohen, J. D. (1997). Dissociating working memory from task difficulty in human prefrontal cortex. Neuropsychologia, 35(10), 1373–1380.

Barch, D. M., Carter, C. S., Braver, T. S. ans Sabb, F. W., MacDonald, A. r., Noll, D. C., & Cohen, J. D. (2001). Selective deficits in prefrontal cortex function in medication-naive patients with schizophrenia. Archives of General Psychiatry, 58, 280–8.

Barkley, R. A. (1997). Behavioral inhibition, sustained attention, and executive functions: Constructing a unifying theory of ADHD. Psychological Bulletin, 121, 65–94.

Barrett, S. L., Bell, R., Watson, D., & King, D. J. (2004). Effects of amisulpiride, risperidone and chlorpromazine on auditory and visual latent inhibition, prepulse inhibition, executive function and eye movements in healthy volunteers. Journal of Psychopharmacology, 18, 156–72.

Baunez, C., & Robbins, T. W. (1999). Effects of dopamine depletion of the dorsal striatum and further interaction with subthalamic nucleus lesions in an attentional task in the rat. Neuroscience, 92, 1343.

Bear, M. F., & Malenka, R. C. (1994). Synaptic plasticity: LTP and LTD. Current Opinion in Neurobiology, 4, 389–399.

Bechara, A., Damasio, H., Tranel, D., & Anderson, S. W. (1998). Dissociation of working memory from decision making within the human prefrontal cortex. Journal of Neuroscience, 18, 428–37.

Beiser, D. G., & Houk, J. C. (1998). Model of cortical-basal ganglionic processing: encoding the serial order of sensory events. Journal of Neurophysiology, 79, 3168–3188.

Ben-Jonathan, N. (1985). Dopamine: a prolactin-inhibiting hormone. Endocrinological Review, 6, 564–89.

Berns, G. S., & Sejnowski, T. J. (1998). A computational model of how the basal ganglia produces sequences. Journal of Cognitive Neuroscience, 10, 108–121.

Berridge, K. C., & Robinson, T. E. (1998). What is the role of dopamine in reward: hedonic impact, reward learning or incentive salience. Brain Research Review, 28, 309–69.

Biederman, J., & Faraone, S. V. (2002). Current concepts on the neurobiology of attention-deficit/hyperactivity disorder. Journal of Attention Disorders, 6 Suppl 1, S7–16.

Biller, B. M., Molitch, M. E., Vance, M. L., Cannistraro, K. B., Davis, K. R., Simons, J., Schoenfelder, J. R., & Klibanski, A. (1996). Treatment of prolactin-secreting macroadenomas with the once-weekly dopamine agonist cabergoline. Journal of Clinical Endicronology & Metabolism, 81, 2338–43.

Black, K. J., Gado, M. H., & Perlmutter, J. S. (1997). PET measurement of dopamine D2 receptor-mediated changes in striatopallidal function. Journal of Neuroscience, 17, 3168–77.

Blackwell, K. T., Czubayko, U., & Plenz, D. (2003). Quantitative estimate of synaptic inputs to striatal neurons during up and down states in vitro. Journal of Neuroscience, 23, 9123–9145.

Blair, J. R., & Spreen, O. (1989). Predicting premorbid IQ: A revision of the National Adult Reading Test. Clinical Neuropsychologist, 3, 129–36.

Bloch, B., & LeMoine, C. (1994). Neostriatal dopamine receptors. Trends in Neurosciences, 17, 3–4.

Blond, G. S., Crepel, F., & Otani, S. (2002). Long-term potentiation in rat prefrontal slices facilitated by phased application of dopamine. European Journal of Pharmacology, 438, 115–116.

Bolla, K. I., Eldreth, D. A., London, E. D., Kiehl, K. A., Mouratidis, M., & Contoreggi, C. e. a. (2003). Orbitofrontal cortex dysfunction in abstinent cocaine abusers performing a decision-making task. Neuroimage, 19, 1085–94.

Braver, T. S., Barch, D. M., & Cohen, J. D. (1999). Cognition and control in schizophrenia: A computational model of dopamine and prefrontal function. Biological Psychiatry, 46, 312–328.

Braver, T. S., Barch, D. M., Keys, B. A., Carter, C. S., Cohen, J. D., Kaye, J. A., Janowsky, J. S., Taylor, S. F., Yesavage, J. A., & Mumenthaler, M. S. (2001). Context processing in older adults: Evidence for a theory relating cognitive control to neurobiology in healthy aging. Journal of Experimental Psychology General, 130, 746–763.

Braver, T. S., & Cohen, J. D. (2000). On the control of control: The role of dopamine in regulating prefrontal function and working memory. In S. Monsell, & J. Driver (Eds.), Control of cognitive processes: Attention and performance XVIII (pp. 713–737). Cambridge, MA: MIT Press.

Braver, T. S., & Cohen, J. D. (2001). Working memory, cognitive control, and the prefrontal cortex: Computational and empirical studies. Cognitive Processing, 2, 25–55.

Brown, J. W., Bullock, D., & Grossberg, S. (2004a). How laminar frontal cortex and basal ganglia circuits interact to control planned and reactive saccades. Neural Networks, 17, 471–510.

Brown, J. W., Bullock, D., & Grossberg, S. (2004b). How laminar frontal cortex and basal ganglia circuits interact to control planned and reactive saccades. Neural Networks, 17, 471–510.

Brown, L. L., Schneider, J. S., & Lidsky, T. I. (1997). Sensory and cognitive functions of the basal ganglia. Current Opinion in Neurobiology, 7, 157.

Brown, R. G., & Marsden, C. D. (1990). Cognitive function in parkinson's disease: From description to theory. Trends in Neurosciences, 13, 21–29.

Brozoski, T. J., Brown, R. M., Rosvold, H. E., & Goldman, P. S. (1979). Cognitive deficit caused by regional depletion of dopamine in prefrontal cortex of rhesus monkey. Science, 205, 929–932.

Bullock, D., & Grossberg, S. (1988). Neural dynamics of planned arm movements: Emergent invariants and speed-accuracy properties during trajectory formation. Psychological Review, 95, 49–90.

Bymaster, F., Perry, K. W., Nelson, D. L., Wong, D. T., Rasmussen, K., Moore, N. A., & Calligaro, D. O. (1999). Olanzapine: a basic science update. British Journal of Psychiatry Supplement, 37, 36–40.

Calabresi, P., Saiardi, A., Pisani, A., Baik, J. H., Centonze, D., Mercuri, N. B., Bernardi, G., & Borrelli, E. (1997). Abnormal synaptic plasticity in the striatum of mice lacking dopamine D2 receptors. Journal of Neuroscience, 17, 4536–44.

Camperi, M., & Manias, M. V. (2003). On duration and dopamine modulation of sustained activity in prefrontal cortex using conductance-based network models. Neurocomputing, 52-54, 699–705.

Camps, M., Cortes, R., Gueye, B., Probst, A., & Palacios, J. M. (1989). Dopamine receptors in the human brain: autoradiographic distribution of D_2 sites. Neuroscience, 28, 275–290.

Cardinal, R. N., Robbins, T. W., & Everitt, B. J. (2000). The effects of d-amphetamine, chlordiazepoxide, alpha-flupenthixol and behavioural manipulations on choice of signalled and unsignalled delayed reinforcement in rats. Psychopharmacology, 152, 362–75.

Carver, A. C., Livesey, D. J., & Charles, M. (2001). Age related changes in inhibitory control as measured by stop signal task performance. International Journal of Neuroscience, 107, 43–61.

Casey, B. J. (2001). Develoment and disruption of inhibitory mechanisms of attention. In R. S. Siegler, & J. L. McClelland (Eds.), Mechanisms of cognitive development: The Carnegie Symposium on Cognition, Vol. 28. Hillsdale, NJ: Erlbaum.

Casey, B. J., Durston, S., & Fosella, J. A. (2001a). Evidence for a mechanistic model of cognitive control. Clinical Neuroscience Research, 1, 267–282.

Casey, B. J., Durston, S., & Fossella, J. A. (2001b). Evidence for a mechanistic model of cognitive control. Clinical Neuroscience Research, 1, 267–282.

Cavada, C., Company, T., Tejedor, J., Cruz-Rizzolo, R. J., & Reinoso-Suarez, F. (2000). The anatomical connections of the macaque monkey orbitofrontal cortex: a review. Cerebral Cortex, 10, 220–42.

Centonze, D., Picconi, B., Gubellini, P., Bernardi, G., & Calabresi, P. (2001). Dopaminergic control of synaptic plasticity in the dorsal striatum. European Journal of Neuroscience, 13, 1071–1077.

Chevalier, G., & Deniau, J. M. (1990). Disinhibition as a basic process in the expression of striatal functions. Trends in Neurosciences, 13, 277–280.

Coffey, S. F., Gudleski, G. D., Saladin, M. E., & Brady, K. T. (2003). Impulsivity and rapid discounting of delayed hypothetical rewards in cocaine-dependent individuals. Experimental and Clinical Psychopharmacology, 11, 18–25.

Cohen, J. D., Braver, T. S., & Brown, J. W. (2002). Computational perspectives on dopamine function in prefrontal cortex. Current Opinion in Neurobiology, 12, 223–229.

Cohen, J. D., Braver, T. S., & O'Reilly, R. C. (1996). A computational approach to prefrontal cortex, cognitive control, and schizophrenia: Recent developments and current challenges. Philosophical Transactions of the Royal Society (London) B, 351, 1515–1527.

Cohen, J. D., Dunbar, K., & McClelland, J. L. (1990). On the control of automatic processes: A parallel distributed processing model of the Stroop effect. Psychological Review, 97(3), 332–361.

Cohen, J. D., Perlstein, W. M., Braver, T. S., Nystrom, L. E., Noll, D. C., Jonides, J., & Smith, E. E. (1997). Temporal dynamics of brain activity during a working memory task. Nature, 386, 604–608.

Colao, A., Lombardi, G., & Annunziato, L. (2000). Cabergoline. Expert Opinion on Pharmacotherapy, 1, 555–74.

Collins, P., Wilkinson, L. S., Everitt, B. J., Robbins, T. W., & Roberts, A. C. (2000). The effect of dopamine depletion from the caudate nucleus of the common marmoset (callithrix jacchus) on tests of prefrontal cognitive function. Behavioral Neuroscience, 114, 3–1–7.

Comings, D. E., Wu, S., Chiu, C., Ring, R. H., Gade, R., & Ahn, C. e. a. (1996). Polygenic inheritance of Tourette syndrome, stuttering, attention deficit hyperactivity, conduct, and oppositional defiant disorder: the additive and subtractive effect fo the three dopaminergic genes–DRD2, D beta H, and DAT1. American Journal of Medical Genetics, 67, 264–88.

Cook, E. H., Stein, M. A., Krasowski, M. D., Cox, N. J., Olkon, D. M., & Kieffer, J. E. (1995). Association of attention deficit disorder and the dopamine transporter gene. American Journal of Human Genetics, 56, 993–998.

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

Cools, R., Barker, R. A., Sahakian, B. J., & Robbins, T. W. (2003). L-Dopa medication remediates cognitive inflexibility, but increases impulsivity in patients with Parkinson's disease. Neuropsychologia, 41, 1431–1441.

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

Corsello, S. M., Ubertini, G., Altomare, M., Lovicu, R. M., Migneco, M. G., Rota, C. A., & Colosimo, C. (2003). Giant prolactinomas in men: efficacy of cabergoline treatment. Clinical Endocrinology, 58, 662–70.

Cowan, R. L., Sesack, S. R., Van Bockstaele, E. J., Branchereau, P., Chan, J., & Pickel, V. M. (1994). Analysis of synaptic inputs and targets of physiologically characterized neurons in rat frontal cortex: combined in vivo intracellular recording and immunolabeling. Synapse, 17, 101–14.

Crofts, H. S., Dalley, J. W., Collins, P., Van Denderen, J. C., Everitt, B. J., Robbins, T. W., & Roberts, A. C. (2001). Differential effects of 6-OHDA lesions of the frontal cortex and caudate nucleus on the ability to acquire an attentional set. Cerebral Cortex, 11, 1015–1026.

Crutcher, M. D., & Alexander, G. E. (1990). Movement-related neuronal activity selectively coding either direction or muscle pattern in three motor areas of the monkey. Journal of Neurophysiology, 64, 151–163.

Daneman, M., & Carpenter, P. A. (1980). Individual differences in working memory and reading. Journal of Verbal Learning and Verbal Behavior, 19, 450–466.

Darby, J. K., Pasta, D. J., Dabiri, L., & Mosbacher, D. (1995). Haloperidol dose and blood level variability: toxicity and interindividual variability in the nonresponder patient in the clinical practice setting. Journal of Clinical Psychopharmacology, 15, 334–40.

de Wit, H., Enggasser, J. L., & Richards, J. B. (2002). Acute administration of d-amphetamine decreases impulsivity in healthy volunteers. Neuropsychopharmacology, 27, 813–25.

Dehaene, S., & Changeux, J. P. (1989). A simple model of prefrontal cortex function in delayed-response tasks. Journal of Cognitive Neuroscience, 1, 244–261.

Delgado, M. R., Locke, H. M., Stenger, V. A., & Fiez, J. A. (2003). Dorsal striatum responses to reward and punishment: Effects of valence and magnitude manipulations. Cognitive, Affective, and Behavioral Neuroscience, 3, 27–38.

Delius, J. D., & Siemann, M. (1998). Transitive responding in animals and humans: Exaptation rather than adaptation? Behavioural Processes, 42, 107–37.

Dellu-Hagedorn, F. (in press). Spontaneous individual differences in cognitive performances of young adult rats predict locomotor response to amphetamine. Neurobiology of Learning and Memory.

DeLong, M. R. (1990). Primate models of movement disorders of basal ganglia origin. Trends in Neurosciences, 13, 281–5.

Depue, R. A., & Collins, P. F. (1999). Neurobiology of the structure of personality: Dopamine, facilitation of incentive motivation, and extraversion. the Behavioral and Brain Sciences, 22, 491.

Diamond, A. (1990). The development and neural bases of memory functions as indexed by the A-not-B task: Evidence for dependence on dorsolateral prefrontal cortex. In A. Diamond (Ed.), The development and neural bases of higher cognitive functions (pp. 267–317). New York: New York Academy of Science Press.

Diamond, A., & Goldman-Rakic, P. S. (1986). Comparative development in human infants and infant rhesus monkeys of cognitive functions that depend on prefrontal cortex. Society for Neuroscience Abstracts, 12, 742.

Diamond, A., & Goldman-Rakic, P. S. (1989). Comparison of human infants and rhesus monkeys on Piaget's A$\overline{B}$ task: Evidence for dependence on dorsolateral prefrontal cortex. Experimental Brain Research, 74, 24–40.

Diamond, A., & Taylor, C. (1996). Development of an aspect of executive control: Development of the abilities to remember what I said and to "Do as I say, not as I do". Developmental Psychology, 29, 315–34.

Dias, R., Robbins, T. W., & Roberts, A. C. (1996). Dissociation in prefrontal cortex of affective and attentional shifts. Nature, 380, 69.

Dias, R., Robbins, T. W., & Roberts, A. C. (1997). Dissociable forms of inhibitory control within prefrontal cortex with an analog of the Wisconsin Card Sort Test: Restriction to novel situations and independence from "on-line" processing. Journal of Neuroscience, 17, 9285–9297.

Dilmore, J. G., Gutkin, B. G., & Ermentrout, G. B. (1999). Effects of dopaminergic modulation of persistent sodium currents on the excitability of prefrontal cortical neurons: A computational study. Neurocomputing, 26, 104–116.

Dominey, P. F. (1995). Complex sensory-motor sequence learning based on recurrent state representation and reinforcement learning. Biological Cybernetics, 73, 265–274.

Dominey, P. F., & Arbib, M. A. (1992). Cortico-subcortical model for generation of spatially accurate sequential saccades. Cerebral Cortex, 2, 153–175.

Dougherty, D. D., Bonab, A. A., Spencer, T. J., Rauch, S. L., Madras, B. K., & Fischman, A. J. (1999). Dopamine transporter density in patients with attention deficit hyperactivity disorder. The Lancet, 354, 2132–2133.

Douglas, R. J., & Martin, K. A. C. (1990). Neocortex. In G. M. Shepherd (Ed.), The synaptic organization of the brain (Chap. 12, pp. 389–438). Oxford: Oxford University Press.

Doya, K. (2000). Complementary roles of the basal ganglia and cerebellum in learning and motor control. Current Opinion in Neurobiology, 10, 732–739.

Dreher, J. C., & Burnod, Y. (2002). An integrative theory of the phasic and tonic modes of dopamine modulation in the prefrontal cortex. Neural Networks, 15, 583–602.

Dubois, B., Malapani, C., Verin, M., Rogelet, P., Deweer, B., & Pillon, B. (1994). Cognitive functions in the basal ganglia: the model of Parkinson disease. Revue Neurologique (Paris), 150, 763–770.

Durstewitz, D., Kelc, M., & Gunturkun, O. (1999). A neurocomputational theory of the dopaminergic modulation of working memory functions. Journal of Neuroscience, 19, 2807.

Durstewitz, D., & Seamans, J. K. (2002). The computational role of dopamine d1 receptors in working memory. Neural Networks, 15, 561–572.

Durstewitz, D., Seamans, J. K., & Sejnowski, T. J. (2000a). Dopamine-mediated stabilization of delay-period activity in a network model of prefrontal cortex. Journal of Neurophysiology, 83, 1733–1750.

Durstewitz, D., Seamans, J. K., & Sejnowski, T. J. (2000b). Neurocomputational models of working memory. Nature Neuroscience, 3 supp, 1184–1191.

Durston, S., Tottenham, N. T., Thomas, K. M., Davidson, M. C., Eigsti, I.-M., Yang, Y., Ulug, A. M., & Casey, B. J. (2003). Differential patterns of striatal activation in young children with and without ADHD. Biological Psychiatry, 53, 871–878.

Elliott, R., Dolan, R. J., & Frith, C. D. (2000). Dissociable functions in the medial and lateral orbitofrontal cortex: evidence from human neuroimaging studies. Cerebral Cortex, 10, 308–17.

Erickson, S. L., & Lewis, D. A. (2000). Prefrontal cortical inputs to monkey mediodorsal thalamus. Society for Neuroscience Abstracts (p. 461). San Diego, CA: Society for Neuroscience.

Erickson, S. L., & Lewis, D. A. (2004). Cortical connections of the lateral mediodorsal thalamus in cynomolgus monkeys. Journal of Comparative Neurology, 473, 107–27.

Eyny, Y. S., & Horvitz, J. C. (2003). Opposing roles of D1 and D2 receptors in appetitive conditioning. Journal of Neuroscience, 23, 1584–1587.

Fellous, J. M., Wang, X. J., & Lisman, J. E. (1998). A role for NMDA-receptor channels in working memory. Nature Neuroscience, 1, 273–275.

Filion, M., & Tremblay, L. (1991). Abnormal spontaneous activity of globus pallidus neurons in monkey with MPTP-induced parkinsonism. Brain Research, 547, 142–151.

Fillmore, M. T., & Rush, C. R. (2002). Impaired inhibitory control of behavior in chronic cocaine users. Drug and Alcohol Dependence, 66, 265–73.

Fillmore, M. T., Rush, C. R., & Hays, L. (2002). Acute effects of oral cocaine on inhibitory control of behavior in humans. Drug and Alcohol Dependence, 67, 157–67.

Finch, D. M. (1999). Plasticity of responses to synaptic inputs in rat ventral striatal neurons after repeated administration of the dopamine D2 antagonist raclopride. Synapse, 31, 297–301.

Fiorillo, C. D., Tobler, P. N., & Schultz, W. (2003). Discrete coding of reward probability and uncertainty by dopamine neurons. Science, 299, 1898–1901.

Fog, R. (1972). On stereotypy and catalepsy: studies on the effect of amphetamines and neuroleptics in rats. Acta Neurol Scand Suppl, 50, 3–66.

Foote, S. L., & Morrison, J. H. (1987). Extrathalamic modulation of cortex. Annual Review of Neuroscience, 10, 67–85.

Fox, C. A., & Rafols, J. A. (1976). The striatal efferents in the globus pallidus and in the substantia nigral. In M. D. Yahr (Ed.), The basal ganglia (pp. 37–55). New York: Raven Press.

Frank, M. J. (in press). Dynamic dopamine modulation in the basal ganglia: A neurocomputational account of cognitive deficits in medicated and non-medicated Parkinsonism. Journal of Cognitive Neuroscience.

Frank, M. J. (under revisions). Dynamic dopamine modulation of striato-cortical circuits in cognition: An integrative account spanning multiple levels of analysis. Behavioral and Brain Sciences.

Frank, M. J., & Claus, E. D. (in preparation). Anatomy of a decision: Striato-orbitofrontal interactions in reinforcement learning, decision making and reversal.

Frank, M. J., Loughry, B., & O'Reilly, R. C. (2001). Interactions between the frontal cortex and basal ganglia in working memory: A computational model. Cognitive, Affective, and Behavioral Neuroscience, 1, 137–160.

Frank, M. J., & O'Reilly, R. C. (in press). A mechanistic account of striatal dopamine function in human cognition: Psychopharmacological studies with cabergoline and haloperidol. Behavioral Neuroscience.

Frank, M. J., Rudy, J. W., Levy, W. B., & O'Reilly, R. C. (in press). When logic fails: Implicit transitive inference in humans. Memory and Cognition.

Frank, M. J., Rudy, J. W., & O'Reilly, R. C. (2003). Transitivity, flexibility, conjunctive representations and the hippocampus: II. A computational analysis. Hippocampus, 13, 341–54.

Frank, M. J., Seeberger, L. C., & O'Reilly, R. C. (2004). By carrot or by stick: Cognitive reinforcement learning in Parkinsonism. Science, 306, 1940–3.

Funahashi, S., Bruce, C. J., & Goldman-Rakic, P. S. (1989). Mnemonic coding of visual space in the monkey's dorsolateral prefrontal cortex. Journal of Neurophysiology, 61, 331–349.

Fuster, J. M. (1997). The prefrontal cortex: Anatomy, physiology and neuropsychology of the frontal lobe, 3rd edition. New York: Lippincott-Raven.

Fuster, J. M. (2001). The prefrontal cortex-an update: Time is of the essence. Neuron, 30, 319–333.

Fuster, J. M., & Alexander, G. E. (1971). Neuron activity related to short-term memory. Science, 173, 652–654.

Garris, P. A., Budygin, E. A., Phillips, P. E., Venton, B. J., Robinson, D. L., Bergstrom, B. P., Rebec, G. V., & Wightman, R. M. (2003). A role for presynaptic mechanisms in the actions of nomifensine and haloperidol. Neuroscience, 118, 819–29.

Gelfand, J., Gullapalli, V., Johnson, M., Raye, C., & Henderson, J. (1997). The dynamics of prefrontal cortico-thalamo-basal ganglionic loops and short term memory interference phenomena. Proceedings of the 19th Annual Conference of the Cognitive Science Society (pp. 253–258). Mahwah, NJ: Erlbaum.

Gerdeman, G. L., Partridge, J. G., Lupica, C. R., & Lovinger, D. M. (2003). It could be habit forming: drugs of abuse and striatal synaptic plasticity. Trends in Neurosciences, 26, 184–192.

Gerfen, C. R. (1992). The neostriatal mosaic: multiple levels of compartmental organization in the basal ganglia. Annual Review of Neuroscience, 15, 285–320.

Gerfen, C. R. (2000). Molecular effects of dopamine on striatal projection pathways. Trends in Neurosciences, 23, S64–S70.

Gerfen, C. R. (2003). D1 dopamine receptor supersensitivity in the dopamine-depleted striatum animal model of Parkinson's disease. Neuroscientist, 9, 455–62.

Gerfen, C. R., & Keefe, K. A. (1994). Neostriatal dopamine receptors. Trends in Neurosciences, 17, 2–3.

Gerfen, C. R., Keefe, K. A., & Gauda, E. B. (1995). D_1 and D_2 dopamine receptor function in the striatum: Coactivation of D_1- and D_2-dopamine receptors on separate populations of neurons results in potentiated immediate early gene response in D_1-containing neurons. Journal of Neuroscience, 15, 8167–8176.

Gerfen, C. R., & Wilson, C. (1996). The basal ganglia. In L. Swanson, A. Bjorkland, & T. Hokfelt (Eds.), Handbook of chemical neuroanatomy. vol 12: Integrated systems of the CNS. (pp. 371–468). Amsterdam: Elsevier.

Gill, M., Daly, G., Heron, S., Hawl, Z., & Fitzgerald, M. (1997). Confirmation of association between attention deficit hyperactivity disorder and a dopamine transporter polymorphism. Molecular Psychiatry, 2, 311–313.

Giros, B., Jaber, M., Jones, S. R., Wightman, R. M., & Caron, M. G. (1996). Hyperlocomotion and indifference to cocaine and amphetamine in mice lacking the dopamine transporter. Nature, 379, 606–612.

Gluck, M. A., Shohamy, D., & Myers, C. (2002). How do people solve the "weather prediction" task?: Individual variability in strategies for probabilistic category learning. Learning and Memory, 9, 408–418.

Gobbel, J. R. (1997). The role of the neostriatum in the execution of action sequences. PhD thesis, University of California, San Diego, San Diego, CA, USA.

Goldman-Rakic, P. (1996). Regional and cellular fractionation of working memory. Proceedings of the National Academy of Sciences, 93, 13473–13480.

Goldman-Rakic, P. S. (1987). Circuitry of primate prefrontal cortex and regulation of behavior by representational memory. Handbook of Physiology — The Nervous System, 5, 373–417.

Goldman-Rakic, P. S. (1995). Architecture of the prefrontal cortex and the central executive. Annals of the New York Academy of Sciences, 769, 71–83.

Goldman-Rakic, P. S., & Brown, R. M. (1982). Postnatal development of monoamine content and synthesis in the cerebral cortex of rhesus monkeys. Brain Research, 256, 339–349.

Goldman-Rakic, P. S., & Friedman, H. R. (1991). The circuitry of working memory revealed by anatomy and metabolic imaging. In H. S. Levin, H. M. Eisenberg, & A. L. Benton (Eds.), Frontal lobe function and dysfunction (pp. 72–91). New York: Oxford University Press.

Goldman-Rakic, P. S., Lidow, M. S., & Gallager, D. W. (1990). Overlap of dopaminergic, adrenergic, and serotoninergic receptors and complementarity of their subtypes in primate prefrontal cortex. Journal of Neuroscience, 10, 2125–2138.

Gorelova, N. A., & Yang, C. R. (2000). Dopamine d1/d5 receptor activation modulates a persistent sodium current in rats prefrontal cortical neurons in vitro. Journal of Neurophysiology, 84, 75.

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

Gottfried, J. A., ODoherty, J., & Dolan, R. J. (2003). Encoding predictive reward value in human amygdala and orbitofrontal cortex. Science, 301, 1104–1106.

Grace, A. A. (1995). The tonic/phasic model of dopamine system regulation: its relevance for understanding how stimulant abuse can alter basal ganglia function. Drug and Alcohol Dependence, 37, 111–29.

Grace, A. A. (2001). Psychostimulant actions on dopamine and limbic system function: relevance to the pathophysiology and treatment of ADHD. In M. V. Solanto, & A. F. Arnsten (Eds.), Stimulant drugs and ADHD: Basic and clinical neuroscience. (pp. 134–55). New York: Oxford University Press.

Grace, A. A., & Onn, S. P. (1989). Morphology and electrophysiological properties of immunocytochemically identified rat dopamine neurons recorded in vitro. Journal of Neuroscience, 9, 3463–3481.

Granon, S., Passetti, F., Thomas, K. L., Dalley, J. W., Everitt, B. J., & Robbins, T. W. (2000). Enhanced and impaired attentional performance after infusion of D1 dopaminergic receptor agents into rat prefrontal cortex. Journal of Neuroscience, 20, 1208–15.

Grant, S., Contoreggi, C., & London, E. D. (2000). Drug abusers show impaired performance in a laboratory test of decision making. Neuropsychologia, 38, 1180–7.

Gray, N. S., Pickering, A. D., Hemsley, D. R., Dawling, S., & Gray, J. A. (1992). Abolition of latent inhibition by a single 5 mg dose of d-amphetamine in man. Psychopharmacology, 107, 425–30.

Graybiel, A. M., & Kimura, M. (1995). Adaptive neural networks in the basal ganglia. In J. C. Houk, J. L. Davis, & D. G. Beiser (Eds.), Models of information processing in the basal ganglia (pp. 103–116). Cambridge, MA: MIT Press.

Groenewegen, H. J., Vermeulen-van Der Zee, E., Te Kortschot, A., & Witter, M. P. (1987). Organization of the projections from the subiculum to the ventral striatum in the rat. a study using anterograde transport of phaseolus vulgaris leucoagglutinin. Neuroscience, 23, 103–20.

Groenewegen, H. J., Wright, C. I., Beijer, A. V., & Voorn, P. (1999). Convergence and segregation of ventral striatal inputs and outputs. Annals of the New York Academy of Sciences, 877, 49–63.

Gurden, H., Takita, M., & Jay, T. M. (2000). Essential role of D1 but not D2 receptors in the NMDA receptor-dependent long-term potentiation at hippocampal-prefrontal cortex synapses in vivo. Journal of Neuroscience, 20, RC106.

Gurden, H., Tassin, J. P., & Jay, T. M. (1999). Integrity of the mesocortical dopaminergic system is necessary for complete expression of in vivo hippocampal-prefrontal cortex long-term potentiation. Neuroscience, 94, 1019.

Gurney, K., Prescott, T. J., & Redgrave, P. (2001). A computational model of action selection in the basal ganglia. I. a new functional anatomy. Biological Cybernetics, 84, 401–410.

Haber, S., Kunishio, K., Mizobuchi, M., & Lynd-Balta, E. (1995). The orbital and medial prefrontal circuit through the primate basal ganglia. Journal of Neuroscience, 15, 4851–67.

Hale, S., Bronik, M. D., & Fry, A. F. (1997). Verbal and spatial working memory in school-age children: developmental differences in susceptibility to interference. Developmental Psychology, 33, 364–371.

Hay, J. F., Moscovitch, M., & Levine, B. (2002). Dissociating habit and recollection: evidence from Parkinson's disease, amnesia and focal lesion patients. Neuropsychologia, 40, 1324–1334.

Hebb, D. O. (1949). The organization of behavior. New York: Wiley.

Henik, A., Singh, J., Beckley, D. J., & Rafal, R. D. (1993). Disinhibtion of automatic word reading in Parkinson's disease. Cortex, 29, 589–99.

Hernandez-Lopez, S., Bargas, J., Surmeier, D. J., Reyes, A., & Galarraga, E. (1997). D1 receptor activation enhances evoked discharge in neostriatal medium spiny neurons by modulating an L-type Ca^{2+} conductance. Journal of Neuroscience, 17, 3334–42.

Hernandez-Lopez, S., Tkatch, T., Perez-Garci, E., Galarraga, E., Bargas, J., Hamm, H., & Surmeier, D. J. (2000). D2 dopamine receptors in striatal medium spiny neurons reduce L-type Ca^{2+} currents and excitability via a novel PLCβ1-IP_3-calcineurin-signaling cascade. Journal of Neuroscience, 20, 8987–95.

Hikosaka, K., & Watanabe, M. (2000). Delay activity of orbital and lateral prefrontal neurons of the monkey varying with different rewards. Cerebral Cortex, 10, 263–71.

Hikosaka, O. (1989). Role of basal ganglia in initiation of voluntary movements. In M. A. Arbib, & S. Amari (Eds.), Dynamic interactions in neural networks: Models and data (pp. 153–167). Berlin: Springer-Verlag.

Hikosaka, O. (1998). Neural systems for control of voluntary action–a hypothesis. Advances in Biophysics, 35, 81–102.

Hochreiter, S., & Schmidhuber, J. (1997). Long short term memory. Neural Computation, 9, 1735–1780.

Hollerman, J. R., & Schultz, W. (1998). Dopamine neurons report an error in the temporal prediction of reward during learning. Nature Neuroscience, 1, 304–309.

Holroyd, C. B., & Coles, M. G. H. (2002). The neural basis of human error processing: Reinforcement learning, dopamine, and the error-related negativity. Psychological Review, 109, 679–709.

Horvitz, J. C. (2000). Mesolimbocortical and nigrostriatal dopamine responses to salient non-reward events. Neuroscience, 96, 651.

Horvitz, J. C. (2001). The effects of D1 and D2 receptor blockade on the acquisition and expression of a conditioned appetitive response. Appetite, 37, 119–20.

Hoshi, E., Shima, K., & Tanji, J. (2000). Neuronal activity in the primate prefrontal cortex in the process of motor selection based on two behavioral rules. Journal of Neurophysiology, 83, 2355.

Houk, J. C., & Wise, S. P. (1995). Distributed modular architectures linking basal ganglia, cerebellum, and cerebral cortex: their role in planning and controlling action. Cerebral Cortex, 5, 95–110.

Hyland, B. I., Reynolds, J. N., Hay, J., Perk, C. G., & Miller, R. (2002). Firing modes of midbrain dopamine cells in the freely moving rat. Neuroscience, 114, 475–492.

Ichikawa, K., & Kojima, M. (2001). Pharmacological effects of cabergoline against parkinsonism. Nippon Yakurigaku Zasshi, 117, 395–400.

Ilgin, N., Senol, S., Gucuyener, K., Gokcora, N., Atavci, S., & Sener, S. (2001). Is increased D2 receptor availability associated with response to stimulant medication in ADHD. Developmental Medicine and Child Neurology, 43, 755–760.

Ince, E., Ciliax, B. J., & Levey, A. I. (1997). Differential expression of D_1 and D_2 dopamine and m4 muscarinic acteylcholine receptor proteins in identified striatonigral neurons. Synapse, 27, 257–366.

Iversen, S., & Mishkin, M. (1970). Perseverative interference in monkeys following selective lesions of the inferior prefrontal convexity. Experimental Brain Research, 11, 376–86.

Jackson, G. M., Jackson, S., Harrison, J., Henderson, L., & Kennard, C. (1995). Serial reaction time learning and Parkinson's disease: evidence for a procedural learning deficit. Neuropsychologia, 33, 577–593.

Jackson, M. E., Frost, A. S., & Moghaddam, B. (2001). Stimulation of prefrontal cortex at physiologically relevant frequencies inhibits dopamine release in the nucleus accumbens. Journal of Neurochemistry, 78, 920–3.

Jackson, M. E., & Moghaddam, B. (2001). Amygdala regulation of nucleus accumbens dopamine output is governed by the prefrontal cortex. Journal of Neuroscience, 21, 676–81.

Jackson, S., & Houghton, G. (1995). Sensorimotor selection and the basal ganglia: A neural network model. In J. C. Houk, J. L. Davis, & D. G. Beiser (Eds.), Models of information processing in the basal ganglia (pp. 337–368). Cambridge, MA: MIT Press.

Jellinger, K. A. (2002). Recent developments in the pathology of Parkinson's disease. Journal of Neural Transmission Supplement, 62, 347–76.

Jiang, H., Stein, B. E., & McHaffie, J. G. (2003). Opposing basal ganglia processes shape midbrain visuomotor activity bilaterally. Nature, 419, 982–985.

Joel, D., & Weiner, I. (2000). The connections of the dopaminergic system with the striatum in rats and primates: an analysis with respect to the functional and compartmental organization of the striatum. Neuroscience, 96, 451.

Jones, G. H., & Robbins, T. W. (1992). Differential effects of mesocortical, mesolimbic and mesostriatal dopamine depletion on spontaneous, conditioned and drug-induced locomotor activity. Pharmacology, Biochemistry and Behavior, 43, 883–895.

Judd, C. M., & McClelland, G. H. (1989). Data analysis, a model-comparison approach. Orlando, FL: Harcourt Brace Jovanovich.

Kaasinen, V., Nurmi, E., Bergman, J., Eskola, O., Solin, O., Sonninen, P., & Rinne, J. O. (2001). Personality traits and brain dopaminergic function in Parkinson's disease. Proceedings of the National Academy of Sciences, 98, 13272–7.

Kakade, S., & Dayan, P. (2002). Acquisition and extinction in autoshaping. Psychological Review, 109, 533–544.

Kapur, S., Remington, G., Jones, C., Wilson, A., DaSilva, J., Houle, S., & Zipursky, R. B. (1996). High levels of dopamine D_2 receptor occupancy with low-dose haloperidol treatment: a PET study. American Journal of Psychiatry, 153, 948–950.

Kaufman, J. N., & Garavan, H. (2004). An fMRI investigation on the effects of intravenous cocaine on cognitive control. Cognitive Neuroscience Society Meeting Poster D46.

Kaufman, J. N., Ross, T. J., Stein, E. A., & Garavan, H. (2003). Cingulate hypoactivity in cocaine users during a GO-NOGO task as revealed by event-related functional magnetic resonance imaging. Journal of Neuroscience, 23, 7839–43.

Kerr, J. N., & Wickens, J. R. (2001). Dopamine D-1/D-5 receptor activation is required for long-term potentiation in the rat neostriatum in vitro. Journal of Neurophysiology, 85, 117–124.

Keshavan, M. S., Sujata, M., Mehra, A., Montrose, D. M., & Sweeney, J. (2003). Psychosis proneness and ADHD in young relatives of schizophrenia patients. Schizophrenia Research, 59, 85–92.

Kimberg, D. Y., & D'Esposito, M. (2003). Cognitive effects of the dopamine receptor agonist pergolide. Neuropsychologia, 41, 1020–1027.

Kimberg, D. Y., D'Esposito, M., & Farah, M. J. (1997). Effects of bromocriptine on human subjects depend on working memory capacity. Neuroreport, 8, 3581–3585.

Kish, S. J., Shannak, K., & Hornykiewicz, O. (1988). Uneven pattern of dopamine loss in the striatum of patients with idiopathic Parkinson's disease. New England Journal of Medecine, 318, 876–880.

Kitai, S. T., Sugimori, M., & Kocsis, J. C. (1976). Excitatory nature of dopamine in the nigro-caudate pathway. Experimental Brain Research, 24, 351–363.

Kiyatkin, E. A., & Rebec, G. V. (1996). Dopaminergic modulation of glutamate-induced excitations of neurons in the neostriatum and nucleus accumbens of awake, unrestrained rats. Journal of Neurophysiology, 75, 142–153.

Klingberg, T., Vaidya, C. J., Gabrieli, J. D., Moseley, M. E., & Hedehus, M. (1999). Myelination and organization of the frontal white matter in children: a diffusion tensor MRI study. Neuroreport, 10, 2817–21.

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

Knowlton, B. J., Squire, L. R., & Gluck, M. A. (1994). Probabilistic category learning in amnesia. Learning and Memory, 1, 1–15.

Krause, K., Dresel, S. H., Krause, J., Kung, H. F., & Tatsch, K. (2000). Increased striatal dopamine transporter in adult patients with attention deficit hyperactivity disroder: effects of methylphenidate as measured by single photon emission computed tomography. Neuroscience Letters, 285, 107–110.

Kroger, J., Nystrom, L., O'Reilly, R. C., Noelle, D., Braver, T. S., & Cohen, J. D. (in preparation). Multiple levels of temporal abstraction in the prefronal cortex: converging results from a computational model and fmri.

Kubota, K., & Niki, H. (1971). Prefrontal cortical unit activity and delayed alternation performance in monkeys. Journal of Neurophysiology, 34, 337–347.

Kumari, V., Corr, P. J., Mulligan, O. F., Cotter, P. A., Checkley, S. A., & Gray, J. A. (1997). Effects of acute administration of d-amphetamine and haloperidol on procedural learning in man. Psychopharmacology, 129, 271–76.

Lange, H., Thorner, G., & Hopf, A. (1976). Morphometric-statistical structure analysis of human striatum, pallidum, and nucleus subthalamicus. III. nucleus subthalamicus. J. Hirnforsh, 17, 31–41.

Le Moine, C., & Bloch, B. (1995). D1 and D2 dopamine receptor gene expression in the rat striatum: sensitive cRNA probes demonstrate prominent segregation of D1 and D2 mRNAs in distinct neuronal populations of the dorsal and ventral striatum. Journal of Comparative Neurology, 355, 418–26.

LeDoux, J. (1998). Fear and the brain: where have we been, and where are we going? Biological Psychiatry, 44, 1229–1238.

LeDoux, J. E. (1992). Brain mechanisms of emotion and emotional learning. Current Opinion in Neurobiology, 2, 191–197.

Levitt, J. B., Lewis, D. A., Yoshioka, T., & Lund, J. S. (1993). Topography of pyramidal neuron intrinsic connections in macaque monkey prefrontal cortex (areas 9 & 46). Journal of Comparative Neurology, 338, 360–376.

Lewis, B. L., & O'Donnell, P. (2000). Ventral tegmental area afferents to the prefrontal cortex maintain membrane potential 'up' states in pyramidal neurons via D1 dopamine receptors. Cerebral Cortex, 10, 1168–1175.

Lidow, M. S., Goldman-Rakic, P. S., Gallager, D. W., & Rakic, P. (1991). Distribution of dopaminergic receptors in the primate cerebral cortex: quantitative autoradiographic analysis using [^{3}H]raclopride, [^{3}H]spiperone and [^{3}H]SCH23390. Journal of Neuroscience, 40, 657–671.

Ljungberg, T., Apicella, P., & Schultz, W. (1992). Responses of monkey dopamine neurons during learning of behavioral reactions. Journal of Neurophysiology, 67, 145–163.

Logan, G. D., Schachar, R. J., & Tannock, R. (2000). Executive control problems in childhood psychopathology: Stop signal studies of attention deficit hyperactivity disorder. In S. Monsell, & J. Driver (Eds.), Control of cognitive processes: Attention and performance XVIII (pp. 653–677). Cambridge, MA: MIT Press.

Luciana, M., Hanson, H. L., & Whitley, C. B. (2004). A preliminary report on dopamine system reactivity in PKU: acute effects of haloperidol on neuropsychological, physiological, and neuroendocrine functions. Psychopharmacology, 175, 18–25.

Maddox, W. T., Ashby, F. G., & Bohil, C. J. (2003). Delayed feedback effects on rule-based and information-integration category learning. Journal of Experimental Psychology Learning Memory and Cognition, 29, 650–662.

Maddox, W. T., & Filoteo, J. V. (2001). Striatal contributions to category learning: Quantitative modeling of simple linear and complex non-linear rule learning in patients with Parkinson's Disease. Journal of the International Neuropsychological Society, 7, 710–727.

Maruyama, W., Naoi, M., & Narabayashi, H. (1996). The metabolism of L-Dopa and L-threo-3,4-dihydroxyphenylserine and their effects on monoamines in the human brain: analysis of the intraventricular fluid from parkinsonian patients. Journal of Neurological Sciences, 139, 141–148.

Mattay, V. S., Callicott, J. H., Bertolino, A., Heaton, I., Frank, J. A., Coppola, R., Berman, K. F., Goldberg, T. E., & Weinberger, D. R. (2000). Effects of dextroamphetamine on cognitive performance and cortical activation. Neuroimage, 12, 268–75.

McAuley, J. H. (2003). The physiological basis of clinical deficits in Parkinson's disease. Progress in Neurobiology, 69, 27–48.

McGowan, S., Lawrence, A. D., Sales, T., Quested, D., & Graby, P. (2004). Presynaptic dopamine dysfunction in schizophrenia: a positron emission tomographic [18f]fluorodopa study. Archives of General Psychiatry, 61, 134–42.

Mehta, M. A., Sahakian, B. J., McKenna, P. J., & Robbins, T. W. (1999). Systemic sulpiride in young adults simulates the profile of cognitive deficits in Parkinson's disease. Psychopharmacology, 146, 162–174.

Mehta, M. A., Swainson, R., Ogilvie, A. D., Sahakian, B. J., & Robbins, T. W. (2001). Improved short-term spatial memory but impaired reversal learning following the dopamine D2 agonist bromocriptine in human volunteers. Psychopharmacology, 159, 10–20.

Mercuri, N., Bernardi, G., Calabresi, P., Cotugno, A., Levi, G., & Stanzione, P. (1985). Dopamine decreases cell excitability in rat striatal neurons by pre- and postsynaptic mechanisms. Brain Research, 358, 110–121.

Middleton, F. A., & Strick, P. L. (2000a). Basal ganglia and cerebellar loops: Motor and cogntive circuits. Brain Research Reviews, 31, 236–250.

Middleton, F. A., & Strick, P. L. (2000b). Basal ganglia output and cognition: Evidence from anatomical, behavioral, and clinical studies. Brain and Cognition, 42, 183–200.

Middleton, F. A., & Strick, P. L. (2002). Basal-ganglia 'projections' to the prefrontal cortex of the primate. Cerebral Cortex, 12, 926–935.

Miller, E. K., Erickson, C. A., & Desimone, R. (1996). Neural mechanisms of visual working memory in prefontal cortex of the macaque. Journal of Neuroscience, 16, 5154.

Mink, J. W. (1996). The basal ganglia: Focused selection and inhibition of competing motor programs. Progress in Neurobiology, 50, 381–425.

Mink, J. W. (2001). Basal ganglia dysfunction in Tourette's syndrome: A new hypothesis. Pediatric Neurology, 25, 190–98.

Miyake, A., & Shah, P. (Eds.). (1999). Models of working memory: Mechanisms of active maintenance and executive control. New York: Cambridge University Press.

Miyashita, Y., & Chang, H. S. (1988). Neuronal correlate of pictorial short-term memory in the primate temporal cortex. Nature, 331, 68–70.

Mobini, S., Body, S., Ho, M. Y., Bradshaw, C. M., Szabadi, E., Deakin, J. F., & Anderson, I. M. (2002). Effects of lesions of the orbitofrontal cortex on sensitivity to delayed and probabilistic reinforcement. Psychopharmacology, 160, 290–8.

Mogenson, G. J., Jones, D. L., & Yim, C. Y. (1980). From motivation to action: functional interface between the limbic system and the motor system. Progress in Neurobiology, 14, 69–87.

Moghaddam, B., & Bunney, B. S. (1990). Acute effects of typical and atypical antipsychotic drugs on the release of dopamine from prefrontal cortex, nucleus accumbens, and stiratum in the rat: an in vivo microdialysis study. Journal of Neurochemistry, 54, 1755–60.

Monchi, O., Petrides, M., Petre, V., Worsley, K., & Dagher, A. (2001). Wisconsin card sorting revisited: Distinct neural circuits participating in different stages of the task identified by event-related functional magnetic resonance imaging. Journal of Neuroscience, 21, 7733–7741.

Monchi, O., Taylor, J. G., & Dagher, A. (2000). A neural model of working memory processes in normal subjects, parkinson's disease and schizophrenia for fMRI design and predictions. Neural Networks, 13, 953.

Montague, P. R., Dayan, P., & Sejnowski, T. J. (1996). A framework for mesencephalic dopamine systems based on predictive Hebbian learning. Journal of Neuroscience, 16, 1936–1947.

Moody, S. L., Wise, S. P., di Pellegrino, G., & Zipser, D. (1998). A model that accounts for activity in primate frontal cortex during a delayed matching-to-sample task. Journal of Neuroscience, 18, 399–410.

Muller, U., von Cramon, D. Y., & Pollman, S. (1998). D1- versus D2-receptor modulation of visuospatial working memory in humans. Journal of Neuroscience, 18, 2720–2728.

Nicola, S. M., Surmeier, J., & Malenka, R. C. (2000). Dopaminergic modulation of neuronal excitability in the striatum and nucleus accumbens. Anuual Review of Neuroscience, 23, 185–215.

Nieoullon, A. (2002). Dopamine and the regulation of cognition and attention. Progress in Neurobiology, 67, 53–83.

Nishi, A., Snyder, G. L., & Greengard, P. (1997). Bidirectional regulation of DARPP-32 phosphorylation by dopamine. Journal of Neuroscience, 17, 8147–8155.

O'Doherty, J., Kringelback, M. L., Rolls, E. T., Hornak, J., & Andrews, C. (2001). Abstract reward and punishment in the human orbitofrontal cortex. Nature Neuroscience, 4, 95–102.

O'Donnell, P., & Grace, A. A. (1995). Synaptic interactions among excitatory afferents to nucleus accumbens neurons: Hippocampal gating of prefrontal cortical input. Journal of Neuroscience, 15, 3622.

O'Reilly, R. C. (1998). Six principles for biologically-based computational models of cortical cognition. Trends in Cognitive Sciences, 2(11), 455–462.

O'Reilly, R. C., Braver, T. S., & Cohen, J. D. (1999). A biologically based computational model of working memory. In A. Miyake, & P. Shah (Eds.), Models of working memory: Mechanisms of active maintenance and executive control. (pp. 375–411). New York: Cambridge University Press.

O'Reilly, R. C., & Frank, M. J. (2003). Making working memory work: A computational model of learning in the frontal cortex and basal ganglia (ICS Tech Report 03-03, revised 8/04). University of Colorado Institute of Cognitive Science.

O'Reilly, R. C., & Frank, M. J. (in press). Making working memory work: A computational model of learning in the frontal cortex and basal ganglia. Neural Computation.

O'Reilly, R. C., & Munakata, Y. (2000). Computational explorations in cognitive neuroscience: Understanding the mind by simulating the brain. Cambridge, MA: MIT Press.

O'Reilly, R. C., Noelle, D., Braver, T. S., & Cohen, J. D. (2002). Prefrontal cortex and dynamic categorization tasks: Representational organization and neuromodulatory control. Cerebral Cortex, 12, 246–257.

O'Reilly, R. C., & Rudy, J. W. (2001). Conjunctive representations in learning and memory: Principles of cortical and hippocampal function. Psychological Review, 108, 311–345.

Packard, M. G., & Knowlton, B. J. (2002). Learning and memory functions of the basal ganglia. Annual Review of Neuroscience, 25, 563–593.

Partiot, A., Verin, M., & Dubois, B. (1996). Delayed response tasks in basal ganglia lesions in man: Further evidence for a striato-frontal cooperation in behavioural adaptation. Neuropsychologia, 34, 709.

Passingham, R. E. (1993). The frontal lobes and voluntary action. Oxford: Oxford University Press.

Peretti, C. S., Danion, J. M., Kauffmann-Muller, F., Grangé, D., Patat, A., & Rosenzweig, P. (1997). Effects of haloperidol and amisulpiride on motor and cognitive skill learning in healthy volunteers. Psychopharmacology, 131, 329–38.

Perlstein, W. M., Dixit, N. K., Carter, C. S., Noll, D. C., & Cohen, J. D. (2003). Prefrontal cortex dysfunction mediates deficits in working memory and prepotent responding in schizophrenia. Biological Psychiatry, 53, 25–38.

Persiani, S., Rocchetti, M., Pacciarini, M. A., Holt, B., Toon, S., & Strolin-Benedetti, M. (1996). The effect of food on cabergoline pharmacokinetics and tolerability in healthy volunteers. Biopharmaceutics and Drug Disposition, 17, 443–55.

Phillips, P. E., Budygin, E. A., Garris, P. A., Robinson, D. L., & Wightman, R. M. (2001). The presynaptic effect of haloperidol delays catalepsy: implications for its neuroleptic action. Neuroscience Abstracts, 27, 2328.

Picconi, B., Centonze, D., Hakansson, K., Bernardi, G., Greengard, P., Fisone, G., Cenci, M. A., & Calabresi, P. (2003). Loss of bidirectional striatal synaptic plasticity in l-dopa-induced dyskinesia. Nature Neuroscience, 6, 501–506.

Poldrack, R. A., Clark, J., PareBlagoev, E. J., Shohamy, D., Moyano, J. C., Myers, C., & Gluck, M. A. (2001). Interactive memory systems in the human brain. Nature, 414, 546–549.

Poldrack, R. A., Prabakharan, V., Seger, C., & Gabrieli, J. D. (1999). Striatal activation during cognitive skill learning. Neuropsychology, 13, 564–574.

Previc, F. H. (1999). Dopamine and the origins of human intelligece. Brain and Cognition, 41, 299–350.

Pucak, M. L., Levitt, J. B., Lund, J. S., & Lewis, D. A. (1996). Patterns of intrinsic and associational circuitry in monkey prefrontal cortex. Journal of Comparative Neurology, 376, 614–630.

Quay, H. C. (1997). Inhibition and attention deficit hyperactivity disorder. Journal of Abnormal Child Psychology, 25, 7–13.

Quinn, N., Illas, A., Lhermitte, F., & Agid, Y. (1981). Bromocriptine in Parkinson's disease: a study of cardiovascular effects. Journal of Neurology, Neurosurgery and Psychiatry, 44, 426–9.

Rao, S. C., Rainer, G., & Miller, E. K. (1997). Integration of what and where in the primate prefrontal cortex. Science, 276, 821–824.

Rao, S. G., Williams, G. V., & Goldman-Rakic, P. S. (1999). Isodirectional tuning of adjacent interneurons and pyramidal cells during working memory: Evidence for microcolumnar organization in PFC. Journal of Neurophysiology, 81, 1903.

Reber, P. J., & Squire, L. R. (1999). Intact learning of artificial grammars and intact category learning by patients with parkinson's disease. Behavioral Neuroscience, 113, 235.

Remy, P., Jackson, P. L., & Ribeiro, M. J. (2000). Relationships between cognitive deficits and dopaminergic function in the striatum of Parkinson's disease patients: a PET study. Neurology, 54, A372.

Reynolds, J. N. J., Hyland, B. I., & Wickens, J. R. (2001). A cellular mechanism of reward-related learning. Nature, 412, 67–69.

Richfield, E. K., Penney, J. B., & Young, A. B. (1989). Anatomical and affinity state comparisons between dopamine D1 and D2 receptors in the rat central nervous system. Neuroscience, 30, 767–77.

Rinne, J. O., Portin, R., Ruottinen, H., Nurmi, E., Bergman, J., Haaparanta, M., & Solin, O. (2000). Cognitive impairment and the brain dopaminergic system in Parkinson's disease: (^{18}F)-fluorodopa positron emission tomographic study. Archives of Neurology, 57, 470–475.

Roberts, A. C., De Salvia, M. A., Wilkinson, L. S., Collins, P., Muir, J. L., Everitt, B. J., & Robbins, T. W. (1994). 6-Hydroxydopamine lesions of the pre-frontal cortex in monkeys enhance performance on an analog of the Wisconsin card sort test: possible interactions with subcortical dopamine. Journal of Neuroscience, 14, 2531–2544.

Robertson, G. S., Vincent, S. R., & Fibiger, H. C. (1992). D_1 and D_2 dopamine receptors differentially regulate c-fos expression in striatonigral and stiratopallidal neurons. Neuroscience, 49, 285–96.

Rogers, R. D., Sahakian, B. J., Hodges, J. R., Polkey, C. E., Kennard, C., & Robbins, T. W. (1998). Dissociating executive mechanisms of task control following frontal damage and Parkinson's disease. Brain, 121, 815–842.

Rolls, E. T. (1996). The orbitofrontal cortex. Philosophical Transactions of the Royal Society of London, B 351, 1433–44.

Rolls, E. T., Thorpe, S. J., Boytim, M., Szabo, I., & Perrett, D. I. (1984). Responses of striatal neurons in the behaving monkey. 3. Effects of intophoretically applied dopamine on normal responsiveness. Neuroscience, 12, 1201–12.

Roman, T., Rohde, L. A., & Hutz, M. H. (2004). Polymorphisms of the dopamine transporter gene: influence on response to methylphenidate in attention deficit-hyperactivity disorder. American Journal of Pharmacogenomics, 4, 83–92.

Rosenberg, D. R., & Lewis, D. A. (1995). Postnatal maturation of the dopaminergic innervation of monkey prefrontal and motor cortices: a tyrosine hydroxylase immunohistochemical analysis. Journal of Comparative Neurology, 358, 383–400.

Rougier, N. P., Noelle, D., Braver, T. S., Cohen, J. D., & O'Reilly, R. C. (submitted). Prefrontal cortex and the flexibility of cognitive control: Rules without symbols.

Rougier, N. P., & O'Reilly, R. C. (2002). Learning representations in a gated prefrontal cortex model of dynamic task switching. Cognitive Science, 26, 503–520.

Rowe, D. C., Stever, C., Giedinghagen, L. N., Gard, J. M., Cleveland, H. H., Terris, S. T., Mohr, J. H., Sherman, S., Abramowitz, A., & Waldman, I. D. (1998). Dopamine DRD4 receptor polymorphism and attention-deficit hyperactivity disorder. Molecular Psychiatry, 3, 419–426.

Sagvolden, T., Johansen, E. B., Aase, H., & Russell, V. A. (2004). A dynamic developmental theory of Attention-Deficit / Hyperactivity Disorder (ADHD) predominantly hyperactive/impulsive and combined subtypes. Behavioral and Brain Sciences, in press.

Salamone, J. D., & Correa, M. (2002). Motivational views of reinforcement: implications for understanding the behavioral functions of nucleus accumbens dopamine. Behavioral Brain Research, 137, 3–25.

Salin, P., Hajji, M. D., & Kerkerian-Le Goff, L. (1996). Bilateral 6-hydroxydopamine-induced lesion of the nigrostriatal dopamine pathway reproduces the effects of unilateral lesion on substance P but not on enkephalin expression in rat basal ganglia. European Journal of Neuroscience, 8, 1746–1757.

Sanberg, P. R. (1980). Haloperidol-induced catalepsy is mediated by postsynaptic dopamine receptors. Nature, 284, 472–3.

Satoh, T., Nakai, S., Sato, T., & Kimura, M. (2003). Correlated coding of motivation and outcome of decision by dopamine neurons. Journal of Neuroscience, 23, 9913–9923.

Sawaguchi, T. (2001). The effects of dopamine and its antagonists on directional delay-period activity of prefrontal neurons in monkeys during an oculomotor delayed-response task. Neuroscience Research, 41, 115–128.

Sawaguchi, T., & Goldman-Rakic, P. S. (1991). D1 dopamine receptors in prefrontal cortex: Involvement in working memory. Science, 251, 947–950.

Schmitz, Y., Benoit-Marand, M., Gonon, F., & Sulzer, D. (2003). Presynaptic regulation of dopaminergic neurotransmission. Journal of Neurochemistry, 87, 273.

Schoemaker, H., Claustre, Y., Fage, D., Rouquier, L., Chergui, K., Curet, O., Oblin, A., Gonon, F., Carter, C., Benavides, J., & Scatton, B. (1997). Neurochemical characteristics of amisulpiride, an atypical dopamine D2/D3 receptor antagonist with both presynaptic and limbic selectivity. Journal of Pharmacology and Experimental Therapeutics, 280, 83–97.

Schoenbaum, G., & Setlow, B. (2001). Integrating orbitofrontal cortex into prefrontal theory: Common processing themes across species and subdivisions. Learning and Memory, 8, 134–147.

Schultz, W. (1998). Predictive reward signal of dopamine neurons. Journal of Neurophysiology, 80, 1.

Schultz, W. (1999). The reward signal of midbrain dopamine neurons. News in Phsyiological Sciences, 14, 249–255.

Schultz, W. (2002). Getting formal with dopamine and reward. Neuron, 36, 241–263.

Schultz, W., Apicella, P., & Ljungberg, T. (1993). Responses of monkey dopamine neurons to reward and conditioned stimuli during successive steps of learning a delayed response task. Journal of Neuroscience, 13, 900–913.

Schultz, W., Apicella, P., Romo, R., & Scarnati, E. (1995). Context-dependent activity in primate striatum reflecting past and future behavioral events. In J. C. Houk, J. L. Davis, & D. G. Beiser (Eds.), Models of information processing in the basal ganglia (pp. 11–28). Cambridge, MA: MIT Press.

Schultz, W., Dayan, P., & Montague, P. R. (1997). A neural substrate of prediction and reward. Science, 275, 1593.

Seamans, J. K., Gorelova, N., Durstewitz, D., & Yang, C. R. (2001). Bidirectional dopamine modulation of gabaergic inhibition in prefrontal cortical pyramidal neurons. Journal of Neuroscience, 21, 3628–3638.

Seamans, J. K., & Yang, C. R. (2004). The principal features and mechanisms of dopamine modulation in the prefrontal cortex. Progress in Neurobiology, 74, 1–57.

Seeman, P. (1987). Dopamine receptors and the dopamine hypothesis of schizophrenia. Synapse, 1, 133–152.

Seeman, P., & Kapur, S. (2000). Schizophrenia: More dopamine, more D2 receptors. Proceedings of the National Academy of Sciences, 97, 7673.

Seeman, P., & Madras, B. (2002). Methylphenidate elevates resting dopamine which lowers the impulse-triggered release of dopamine: a hypothesis. Behavioural Brain Research, 130, 79–83.

Servan-Schreiber, D., Cohen, J. D., & Steingard, S. (1997). Schizophrenic deficits in the processing of context: A test of a theoretical model. Archives of General Psychiatry, 53, 1105–1113.

Servan-Schreiber, D., Printz, H., & Cohen, J. D. (1990). A network model of catecholamine effects: Gain, signal-to-noise ratio, and behavior. Science, 249, 892–895.

Seung, H. S. (1998). Continuous attractors and oculomotor control. Neural Networks, 11, 1253.

Shima, K., & Tanji, J. (1998). Both supplementary and presupplementary motor areas are crucial for the temporal organization of multiple movements. Journal of Neurophysiology, 80, 3247–3260.

Shohamy, D., Myers, C. E., Grossman, S., Sage, J., Gluck, M. A., & Poldrack, R. A. (2004). Cortico-striatal contributions to feedback-based learning: converging data from neuroimaging and neuropsychology. Brain, 127, 851–9.

Siemann, M., & Delius, J. D. (1998). Algebraic learning and neural network models for transitive and non-transitive responding. European Journal of Cognitive Psychology, 10, 307–334.

Smith, A. G., Neill, J. C., & Costall, B. (1999). The dopamine D2/D3 receptor agonist 7-OH-DPAT induces cognitive impairment in the marmoset. Pharmacology, Biochemistry and Behavior, 63(2), 201–211.

Solanto, M. V. (2002). Dopamine dysfunction in AD/HD: integrating clinical and basic neuroscience research. Behavioral Brain Research, 130, 65–71.

Soliveri, P., Brown, R. G., Jahanshahi, M., Caraceni, T., & Marsden, C. D. (1997). Learning manual pursuit tracking skills in patients with Parkinson's disease. Brain, 120, 1325–1337.

Starke, K., Gothert, M., & Kilbinger, H. (1989). Modulation of neurotransmitter release by presynaptic autoreceptors. Physiology Reviews, 69, 864–989.

Stern, C., & Passingham, R. (1995). The nucleus accumbens in monkeys (macaca fascicularis). Experimental Brain Research, 106, 239–247.

Stevenson, C. W., & Gratton, A. (2003). Basolateral amygdala modulation of the nucleus accumbens dopamine response to stress: role of the medial prefrontal cortex. European Journal of Neuroscience, 17, 1287.

Stocchi, F., Vacca, L., Berardelli, A., Onofrj, M., Manfredi, M., & Ruggieri, S. (2003). Dual dopamine agonist treatment in Parkinson's disease. Journal of Neurology, 250, 822–826.

Suri, R. E., & Schultz, W. (1999). A neural network model with dopamine-like reinforcement signal that learns a spatial delayed response task. Neuroscience, 91, 871–890.

Surmeier, D. J., Bargas, J., Hemmings, H. C., Nairn, A. C., & Greengard, P. (1995). Modulation of calcium currents by a D1 dopaminergic protein kinase/phosphotase cascade in rat neostriatal neurons. Neuron, 14, 385–397.

Surmeier, D. J., Song, W. J., & Yan, Z. (1996). Coordinated expression of dopamine receptors in neostriatal medium spiny neurons. Journal of Neuroscience, 16, 6579–91.

Sutton, R. S. (1988). Learning to predict by the method of temporal diferences. Machine Learning, 3, 9–44.

Swainson, R., Rogers, R. D., Sahakian, B. J., Summers, B. A., Polkey, C. E., & Robbins, T. W. (2000). Probabilistic learning and reversal deficits in patients with Parkinson's disease or frontal or temporal lobe lesions: Possible adverse effects of dopaminergic medication. Neuropsychologia, 38, 596.

Szele, F. G., Artymyshyn, R., Molinoff, P. B., & Chesselet, M. F. (1991). Heterogeneous distribution of dopamine D2 receptor mRNA in the rat striatum: a quantitative analysis with in situ hybridization histochemistry. Anatomical Record, 231, 548–558.

Tam, S. Y., Elsworth, J. D., Bradberry, C. W., & Roth, R. H. (1990). Mesocortical dopamine neurons: high basal firing frequency predicts tyrosine dependence of dopamine synthesis. Journal of Neural Transmission, 81, 97–110.

Tannock, R., Schachar, T. J., & Logan, G. D. (1995). Methylphenidate and cognitive flexibility: Dissociated dose effects on behavior and cognition in hyperactive children. Journal of Abnormal Child Psychology, 23, 235–266.

Taylor, J. G., & Taylor, N. R. (2000). Analysis of recurrent cortico-basal ganglia-thalamic loops for working memory. Biological Cybernetics, 82, 415–432.

Taylor, N. R., & Taylor, J. G. (1999). Modelling the frontal lobes in health and disease. Proceedings of the International Conference on Artificial Neural Networks, 401–406.

Thierry, A. M., Tassin, J. P., Blanc, G., Stinus, L., Scatton, B., & Glowinski, J. (1977). Discovery of the mesocortical dopaminergic system: some pharmacological and functional characteristics. Advances in Biochemistry and Psychopharmacology, 16, 5–12.

Thomas-Ollivier, V., Reymann, J. M., Le Moal, S., Schueck, S., Lieury, A., & Allain, H. (1999). Procedural memory in recent-onset Parkinson's disease. Dementia and Geriatric Cognitive Disorders, 10, 172–180.

Tomer, R., & Aharon-Peretz, J. (2004). Novelty seeking and harm avoidance in Parkinson's disease: effects of asymmetric dopamine deficiency. Journal of Neurology, Neurosurgery & Psychiatry, 75, 972–5.

Tremblay, L., & Schultz, W. (1999). Relative reward preference in primate orbitofrontal cortex. Nature, 398, 704.

Tsung, F.-S., & Cottrell, G. W. (1993). Learning simple arithmetic procedures. Connection Science, 5, 37–58.

Ungless, M. A., Magill, P. J., & Bolam, J. P. (2004). Uniform inhibition of dopamine neurons in the ventral tegmental area by aversive stimuli. Science, 303, 2040–2.

Uziello, A., Baik, J. H., Rouge-Pont, F., Picetti, R., Dierich, A., Lemeur, M., Piazza, P. V., & Borrelli, E. (2000). Distinct functions of the two isoforms of dopamine D2 receptors. Nature, 408, 199–203.

VanElzakker, M., O'Reilly, R. C., & Rudy, J. W. (2003). Transitivity, flexibility, conjunctive representations and the hippocampus: I. an empirical analysis. Hippocampus, 13, 334–340.

Vastag, B. (2001). Pay attention: Ritalin acts much like cocaine. The Journal of the American Medical Association, 286, 905–906.

Volkow, N. D., Chang, L., Wang, G. J., Fowler, J. S., Ding, Y. S., Sedler, M., Logan, J., Franceschi, D., Gatley, J., Hitzemann, R., Gifford, A., Wong, C., & Pappas, N. (2001a). Low level of brain dopamine D2 receptors in methamphetamine abusers: association with metabolism in the orbitofrontal cortex. American Journal of Psychiatry, 158, 2015–21.

Volkow, N. D., Ding, Y.-S., Fowler, J. S., Wang, D. J., Logan, J., Gatley, S. J., Dewey, S. L., Ashby, C., Lieberman, J., Hitzemann, R., & Wolf, A. P. (1995). Is methylphenidate like cocaine? Studies on their pharmokinetics and distribution in human brain. Archives of General Psychiatry, 52, 456–463.

Volkow, N. D., Ding, Y. S., Fowler, J. S., Wang, G. J., Logan, J., Gatley, S. J., Hitzemann, R., Smith, G., Fields, S. D., & Gur, R. (1996a). Dopamine transporters decrease with age. Journal of Nuclear Medicine, 37, 554–559.

Volkow, N. D., Fowler, J. S., & Wang, G. J. (1999). Imaging studies on the role of dopamine in cocaine reinforcement and addiction in humans. Journal of Psychopharmacology, 13, 337–45.

Volkow, N. D., Fowler, J. S., & Wang, G. J. (2003). The addicted human brain: insights from imaging studies. Journal of Clinical Investigation, 111, 1444–1451.

Volkow, N. D., Fowler, J. S., Wang, G. J., Ding, Y. S., & Gatley, S. J. (2002a). Role of dopamine in the therapeutic and reinforcing effects of methylphenidate in humans: results from imaging studies. European Neuropsychopharmacology, 12, 557–66.

Volkow, N. D., Fowler, J. S., Wang, G. J., & Gatley, S. J. (2002b). Mechanism of action of methylphenidate: insights from PET imaging studies. Journal of Attention Disorders, 6, S31–S43.

Volkow, N. D., Ruben, C. J., Wang, G. J., Fowler, J. S., Moberg, P. J., Yu-Shin, D., Hitzemann, R., Smith, G., & Logan, J. (1998). Association between decline in brain dopamine activity with age and cognitive and motor impairment in healthy individuals. American Journal of Psychiatry, 155, 344–349.

Volkow, N. D., Wang, G. J., Fowler, J. S., Logan, J., Gatley, S. J., MacGregor, R. R., Schyler, D., Hitzemann, R., & Wolf, A. P. (1996b). Measuring age-related changes in DA D2 receptors with [^{11}C]raclopride and with [^{18}F]N-methylspiroperidol. Psychiatry Research, 67, 11–16.

Volkow, N. D., Wang, G. J., Fowler, J. S., Logan, J., Gerasimov, M., Maynard, L., Ding, Y.-S., Gatley, S. J., Gifford, A., & Franceschi, D. (2001b). Therapeutic doses of methylphenidate significantly increase extracellular dopamine in the human brain. Journal of Neuroscience, 21, RC121.

von Fersen, L., Wynne, C. D. L., Delius, J. D., & Staddon, J. E. R. (1991). Transitive inference in pigeons. Journal of Experimental Psychology: Animal Behavior Processes, 17, 334–341.

Vriezen, E. R., & Moscovitch, M. (1990). Memory for temporal order and conditional associative-learning in patients with Parkinson's disease. Neuropsychologia, 28, 1283–1293.

Wade, T. R., de Wit, H., & Richards, J. B. (2000). Effects of dopaminergic drugs on delayed reward as a measure of impulsive behavior in rats. Psychopharmacology, 150, 90–101.

Waldman, I. D., Rowe, D. C., Abramowitz, A., Kozel, S. T., Mohr, J. H., Sherman, S. L., Cleveland, H. H., Sanders, M. L., Gard, J. M., & Stever, C. (1998). Association and linkage of the dopamine transporter gene and attention-deficit hyperactivity disorder in children: heterogeneity owing to diagnostic subtype and severity. American Journal of Human Genetics, 63, 1767–1776.

Wallis, J. D., & Miller, E. K. (2003). Neuronal activity in primate dorsolateral and orbital prefrontal cortex during performance of a reward preference task. European Journal of Neuroscience, 18, 2069–81.

Wang, M., Vijayraghavan, S., & Goldman-Rakic, P. S. (2004). Selective D2 receptor actions on the functional circuitry of working memory. Science, 303, 853–6.

Wang, X.-J. (1999). Synaptic basis of cortical persistent activity: The importance of NMDA receptors to working memory. Journal of Neuroscience, 19, 9587.

Waszczak, B. L., Martin, L. P., Finlay, H. E., Zahr, N., & Stellar, J. R. (2002). Effects of individual and concurrent stimulation of striatal D_1 and D_2 dopamine receptors on electrophysiological and behavioral output from rat basal ganglia. Pharamacology and Experimental Therapeutics, 300, 850–61.

Watanabe, M. (1992). Frontal units of the monkey coding the associative significance of visual and auditory stimuli. Experimental Brain Research, 89, 233–247.

Watanabe, M., Kodama, T., & Hikosaka, K. (1997). Increase of extracellular dopamine in primate prefrontal cortex during a working memory task. Journal of Neurophysiology, 78, 2795.

Webster, J. e. a. (1994). Comparison of cabergoline and bromocriptine in the treatment of hyperprolactinemic amenorrhea. New England Journal of Medicine, 31, 904–909.

Weinberger, D. R. (1987). Implications of normal brain development for the pathogenesis of schizophrenia. Archives of General Psychiatry, 44, 660–669.

Wichmann, T., & DeLong, M. R. (2003). Pathophysiology of Parkinson's disease: The MPTP primate model of the human disorder. Annals of the New York Academy of Sciences, 991, 199–213.

Wickens, J. (1993). A theory of the striatum. Oxford: Pergamon Press.

Wickens, J. (1997). Basal ganglia: Structure and computations. Network: Computation in Neural Systems, 8, R77–R109.

Wickens, J. R., Kotter, R., & Alexander, M. E. (1995). Effects of local connectivity on striatal function: simulation and analysis of a model. Synapse, 20, 281–298.

Willcutt, E. G., Brodsky, K., Chhabildas, N., Shanahan, M., Yerys, B., Scott, A., & Pennington, B. (in press). The neuropsychology of ADHD: Validity of the executive function hypothesis. In D. Gozal, & D. Molfese (Eds.), Attention deficit hyperactivity disorder: From genes to animal models to patients.

Williams, G. V., & Goldman-Rakic, P. S. (1995). Modulation of memory fields by dopamine d1 receptors in prefrontal cortex. Nature, 376, 572–575.

Wilson, C. J. (1993). The generation of natural firing patterns in neostriatal neurons. In G. W. Arbuthnott, & P. C. Emson (Eds.), Chemical signalling in the basal ganglia (progress in brain research, vol. 99) (pp. 277–297). Amsterdam: Elsevier.

Wilson, C. J., & Kawaguchi, Y. (1996). The origins of two-state spontaneous membrane potential fluctuations of neostriatal spiny neurons. Journal of Neuroscience, 16, 2397–2410.

Wise, R. A. (1985a). The anhedonia hypothesis: Mark III. Behavioral Brain Science, 8, 178–86.

Wise, R. A. (2002). Brain reward circuitry: Insights from unsensed incentives. Neuron, 36, 229–240.

Wise, R. A., & Rompre, P.-P. (1989). Brain dopamine and reward. Annual Review of Psychology, 40, 191–225.

Wise, S. P. (1985b). The primate premotor cortex: Past, present, and preparatory. Annual Review of Neuroscience, 8, 1–19.

Woodward, T. S., Bub, D. N., & Hunter, M. A. (2002). Task switching deficits associated with Parkinson's disease reflect depleted attentional resources. Neuropsychologia, 40, 1948–1955.

Wu, Q., Reith, M. E., Walker, Q. D., Kuhn, C. M., Caroll, F. I., & Garris, P. A. (2002). Concurrent autoreceptor-mediated control of dopamine release and uptake during neurotransmission: an in vivo voltammetric study. Journal of Neuroscience, 22, 6272–81.

Wynne, C. D. (1995). Reinforcement accounts for transitive inference performance. Animal Learning & Behavior, 23, 207–217.

Wynne, C. D. (1998). A minimal model of transitive inference. In C. D. Wynne, & J. E. Staddon (Eds.), Models of action (pp. 269–307). New Jersey: Lawrence Erlbaum Associates.

Yang, C. R., & Seamans, J. K. (1996). Dopamine D1 receptor actions in layer V-VI rat prefrontal cortex neurons in vitro: Modulation of dendritic-somatic signal integration. Journal of Neuroscience, 16, 1922.

You, Z. B., Tzschentke, T. M., Brodin, E., & Wise, R. A. (1998). Electrical stimulation of the prefrontal cortex increases cholecystokinin, glutamate, and dopamine release in the nucleus accumbens: an in vivo microdialysis study in freely moving rats. Journal of Neuroscience, 18, 6492–500.

Young, A. M. J., Ahier, R. G., & Gray, J. A. (1998). Increased extracellular dopamine in the nucleus accumbens of the rat during associative learning of neutral stimuli. Neuroscience, 83, 1175.

Zald, D. H., Boileau, I., El-Dearedy, W., Gunn, R., McGlone, F., Dichter, G. S., & Dagher, A. (2004). Dopamine transmission in the human striatum during monetary reward tasks. Journal of Neuroscience, 24, 4105–12.

Zink, C. F., Pagnoni, G., Martin-Skurski, M. E., C., C. J., & Berns, G. S. (2004). Human striatal responses to monetary reward depend on saliency. Neuron, 42, 509–17.

Zipser, D., Kehoe, B., Littlewort, G., & Fuster, J. (1993). A spiking network model of short-term active memory. Journal of Neuroscience, 13, 3406–3420.

Appendix A

Implementational details for computational model of basal ganglia/dopamine in cognition

The model is implemented using a subset of the Leabra framework (O'Reilly & Munakata, 2000; O'Reilly, 1998). The two relevant properties of this framework for the present model are a) the use of a point neuron activation function; and b) the k-Winners-Take-All (kWTA) inhibition function that models the effects of inhibitory neurons. These two properties are described in detail in the above references, and also in Frank et al. (2001). Only specific methods related to the present model are described below.

A.1 Parameters for D1 Contrast Enhancement

A simplified version of the Leabra activation function is presented here, to provide context for the parameters associated with contrast enhancement.

Activation communicated to other cells (y_j) is a thresholded (Θ) sigmoidal function of the membrane potential with gain parameter γ:

$$y_j(t) = \frac{1}{\left(1 + \frac{1}{\gamma[V_m(t) - \Theta]_+}\right)} \tag{A.1}$$

where $[x]_+$ is a threshold function that returns 0 if $x < 0$ and x if $X > 0$. In actual implementation, a less discontinuous deterministic function with a softer threshold is used (see O'Reilly, 1998; O'Reilly & Munakata, 2000), but the differences do not effect the contrast enhancement manipulations.

The default activation gain, gamma is 600. The default membrane potential firing threshold, Theta, is 0.25. These parameters were used for tonic levels of DA. For contrast enhancement during phasic DA spikes, the activation gain was increased to 10000*k, and the threshold was increased to 0.25 + 0.04*k, where k is the percentage of intact SNc units (k=1 for control networks; k=0.25 for PD networks). This has the effect of suppressing units that do not meet the higher threshold, but enhancing activity in units that are above this threshold. During phasic dips of DA, the activation gain was reduced to 600 - k*300, and the threshold was 0.25.

A.2 Probabilistic Classification Simulation

For the weather prediction task, the same probabilistic structure was used as in the original study (Knowlton et al., 1994), in terms of both frequency of presentation of individual cue combinations, and their probability of being associated with an outcome of "rain" or "sun". Thus, patterns were presented to the network consisting of one to three cues in blocks of 100 trials. Each cue was represented by a single column of units in the input layer. Thus, a trial that includes cues 1 and 3 together was simulated by activating the first and third column of units in the input (figure 2.2). This cue combination was presented in six out of 100 trials (for frequency of 6 %), of which five of them would involve positive feedback for a rain response, and negative feedback for a sun response (for a probability of 83.3 % rain). The two potential responses in premotor cortex were left and right, corresponding to buttons pushed by participants to respond "sun" or "rain".

A.3 Probabilistic Reversal Simulation

Just as in the WP task, each of the two stimuli in the PR task were represented by a column of units in the input layer. Unlike the WP task, the potential responses involve directly selecting one of the two stimuli in the input (i.e. two alternative forced choice). The actual motor response (i.e., left/right) is not as relevant in this case, because the correct stimulus appears just as often on the left and right sides of the screen. Instead, responses are likely selected **relative** to a particular stimulus that is being considered; the participant can either select it, or switch to the other stimulus. To address this in the model, a **stimulus selection process** was implemented. In any given trial, attention is randomly directed to one stimulus with only contextual information about the other. Potential responses were simply to "approach" the attended stimulus, or to "switch" to the context stimulus. This was modeled by making one of the stimuli more salient: the attended stimulus had all five units in its column fully activated, whereas the context stimulus had only 3 (randomly selected) units weakly activated, with a mean activation of 0.25 and a variance of 0.35. A similar method was implemented to model a two alternative forced choice task in previous work (Frank et al., 2003).

Appendix B

Parkinson Cognitive Study: Additional Analyses

B.1 Model Methods & Statistics

As discussed in chapter 4, the observed behavioral pattern of results in medicated and non-medicated PD patients was predicted by our computational model (Frank, in press). To explicitly compare these results to the model, we simulated the probabilistic selection task in with three model manipulations (intact, simulated PD, and simulated DA medication, as described in the main paper). 25 networks with different sets of random initial synaptic weights were run with each manipulation. Following training, Go/NoGo associations were recorded from the model's striatum in response to input stimuli A and B (Figure 2c in the main paper). For positive Go learning, we computed Go - NoGo associations for stimulus A. For negative NoGo learning, we computed NoGo - Go associations for stimulus B.

We then performed a GLM exactly as in the behavioral analyses. This analysis revealed a highly significant interaction between stimulus condition and the contrast between simulated medication and simulated PD ($F[1,72] = 20.1$, $p < .0001$). This is similar to the behavioral interaction observed between choosing A and avoiding B and the ON/OFF medication contrast in PD patients. Networks with simulated DA medication had significantly higher overall Go associations for stimulus A then did networks with simulated PD [$F(1,72) = 13.6$, $p = .0004$]. Conversely, PD networks had greater NoGo associations for stimulus B than did those with simulated medication [$F(1,72) = 9.2$, $p = .003$].

The observed significant differences between simulated PD and medication networks in Go/NoGo learning could result either from enhancements or impairments in the respective conditions (or both). To disentangle these possibilities, we performed separate planned comparisons that compare Go/NoGo associations of each of the DA manipulations networks to those of the intact networks. Compared with intact networks, simulated PD resulted in significantly lower Go associations to stimulus A [$F(1,48) = 7.8$, $p = .0075$], with non-significantly greater NoGo associations to stimulus B [$F(1,48) = 0.93$, n.s.]. Conversely, networks with simulated DA medication had significantly weaker NoGo associations to stimulus B [$F(1,48) = 4.33$, $p = .04$], and numerically but not significantly stronger Go associations to stimulus A [$F(1,48) = 0.75$, n.s.].

B.1.1 Model Interpretation Issues

There might appear to be some conflict between the way the model is described in the Frank (in press) paper (Frank, in press), and the way it is characterized here. Specifically, whereas in this chapter we emphasize the NoGo bias effects of low tonic dopamine levels in Parkinson networks, Frank (in press) also discusses the decreased "dynamic range" effects of decreased dopamine levels on phasic dips. In fact, both effects are present in the model. Thus, the model has an overall NoGo bias, which means that if it is going to learn anything at all, it will be more able to learn to avoid negative stimuli than to choose positive stimuli. However, because the amount of tonic dopamine is already low, phasic dips during negative feedback actually have less of a differential effect on NoGo learning than in intact networks. This dynamic range issue may result in decreased ability to resolve fine differences in Go and NoGo associations across multiple cue configurations (as in more

complex probabilistic categorization tasks (Knowlton et al., 1996)). However, this issue is not relevant in tasks that allow one to adopt either a simple Go or NoGo strategy, without having to resolve fine differences. For instance, in our probabilistic selection task, adopting a global NoGo strategy can lead to preserved training performance because each trial always involves one stimulus that has an overall negative association and should be avoided.

B.2 Post-Experiment Questionnaire Analysis for the Transitive Inference Task

Seven questions were asked, as follows.

(1) Do you have any prior knowledge of the symbols used in the experiment?

(2) If you answered "Yes" to question 1, please indicate to what extent you are familiar with these characters.

(3) Did you have the impression that some of the pairs were easier to choose from than others?

(4) Did you think any of the symbols were ALWAYS correct (no matter what the other symbol was)?

(5) Did you think any of the symbols were ALWAYS incorrect (no matter what the other symbol was)?

(6) Did you have the impression that there was some kind of logical rule, order, or hierarchy of symbols in the experiment? If so, please explain briefly.

(7) In the test phase, were there any new symbols or new combinations of symbols?

(8) If you answered "Yes" to question 7, how did you make your choice in these cases? (e.g., guessed, went with instinct, used some sort of rule - explain)

Awareness judgments were made by assessing the above written questionnaires, and asking participants to clarify some responses, while being completely blind to their performance (Frank et al., in press).

Only one senior was judged to be explicitly aware of the logical hierarchy ordering. The remaining 16 seniors and 17 PD patients were judged to be unaware of any notion of logical order or hierarchy among premise pairs. Eight out of 16 seniors, 8 out of 17 patients ON medication, and 4 out of 8 patients OFF medication did in fact notice that there was a stimulus that was always correct (stimulus A). Nine seniors, 6 patients ON medication, and 4 patients OFF medication noticed that there was a stimulus that was always incorrect (stimulus E). There were no differences between groups in these explicit awareness measures. When asked to describe the "rule", some of these participants stated that they memorized specific pairs, but could not describe any notion of logical order, and didn't explicitly know how to respond to the novel test pairs because they had not memorized the correct response to them during training. Moreover, the "unaware" participants that advanced to the test phase did not use any logical rule or order to determine their choices during test. Many didn't notice there were novel test pairs that differed from the training pairs, and those that did simply "guessed", or went with "instinct." A few participants employed an explicit rule that was incorrect (e.g., "I chose the symbol that was widest").

Appendix C

Psychopharmacological Studies: Drug Effects on Physiological and Subjective Measures

C.1 Drug Effects on Physiological Measures

C.1.1 Drug Effects on Serum Prolactin Secretion

Prolactin levels were obtained before drug ingestion and four hours later, after cognitive tests. There was a main effect of drug on prolactin levels ($F[2,26] = 30.4$, $p < .0001$), such that both active drugs effectively modulated prolactin secretion in opposite directions, as predicted by D2 receptor stimulation/blockade in prolactin-secreting neurons in the pituitary. Although placebo was associated with diminished prolactin levels over this four hour period (prolactin levels normally decrease during the day), cabergoline significantly decreased prolactin secretion beyond that observed under placebo ($F[1,26] = 36.8$, $p < .0001$). In contrast, haloperidol resulted in increased prolactin levels, which was significant relative to placebo ($F[1,26] = 6.3$, $p = .018$).

C.1.2 Drug Effects on Pulse Rate

There was a main effect of drug on pulse rate ($F[2,27] = 4.9$, $p = .015$). Planned contrasts revealed that cabergoline increased pulse rate relative to placebo ($F[1,27] = 9.5$, $p = .0046$), while haloperidol had no significant effect on pulse rate ($F[1,27] = 1.2$, n.s.). The slightly but significantly increased pulse rates by cabergoline are similar to cardiovascular effects caused by other D2 agonists (Quinn, Illas, Lhermitte, & Agid, 1981).

C.2 Drug Effects on Subjective Rating Scales

The mean change scores for the subjective rating scales are shown in Table C.1. There were no significant reported side physical effects between drug conditions, although there was a trend for cabergoline to be associated with increased likelihood of developing a headache ($p = 0.06$). For mental side effects, all drugs were associated with decreases in good mood, but there were no

differences between either of the active drugs and placebo. The only significant within-subject difference was that haloperidol was associated with **less** restlessness compared with placebo ($p = 0.025$). There were trends for haloperidol to be associated with less clearheadedness ($p = .06$) and more drowsiness ($p = 0.16$) compared with placebo. None of these effects were confounding with interpretation of the results reported in the main paper, which involved both enhancements and impairments depending on task condition.

Likert Scale	Placebo	Cabergoline	Haloperidol
Headache	+0.11 (0.1)	+0.46 (0.14)	+0.23 (0.08)
Stomach Ache	0 (0)	+0.04 (0.04)	0 (0)
Nausea	0 (0.05)	-0.04 (0.07)	+0.04 (0.04)
Dizziness	+0.15 (0.11)	+0.25 (0.13)	+0.19 (0.09)
Blurred Vision	+0.11 (0.08)	+0.25 (0.14)	+0.15 (0.1)
Muscle Pain	-0.04 (0.04)	0 (0)	0 (0.06)
Muscle Stiffness	0 (0.05)	+0.13 (0.15)	+0.15 (0.09)
Muscle Twitches	0 (0.05)	0 (0)	+0.15 (0.12)
Good Mood	-0.48 (0.14)	-0.46 (0.18)	-0.38 (0.12)
Bad Mood	+0.19 (0.13)	+0.04 (0.18)	+0.27 (0.2)
Restlessness	+0.59 (0.22)	+0.54 (0.22)	+0.27 (0.17)*
Tenseness	+0.22 (0.11)	0 (0.09)	0.12 (0.16)
Relaxedness	-0.74 (0.17)	-0.67 (0.17)	-0.58 (0.19)
Alertness	-0.07 (0.18)	-0.21 (0.15)	-0.27 (0.18)
Tired/Drowsiness	0.04 (0.18)	0.38 (0.19)	0.54 (0.2)
Clearheadedness	-0.11 (0.2)	-0.29 (0.16)	-0.65 (0.18)
Thought on drug?	2.18 (0.23)	2.42 (0.29)	2.15 (0.23)

Table C.1: Subjective rating scale measures. Participants filled out questionnaires before drug ingestion and after completion of cognitive experiments, using a 5-point Likert scale. Values indicate mean (standard error) change scores from baseline (i.e. score post-tablet – score pre-tablet). The only significant within-subject drug difference was that haloperidol resulted in **less** restlessness compared to the placebo condition. In the post-experimentation questionnaire, participants were asked to rate on a 5-point scale whether they thought they were on an active drug during experimentation. * Significant at the .05 level compared to placebo.