Reinforcement Learning and Semantic Selection: The Role of the Basal Ganglia in Language

A dissertation

submitted in partial fulfillment of the

requirements for the degree of

Doctor of Philosophy

University of Washington

2019

Reading Committee:

Chantel S. Prat, Chair

Andrea Stocco

Lee Osterhout

Geoffrey Boynton

Program Authorized to Offer Degree:

Psychology

University of Washington

Abstract

Reinforcement Learning and Semantic Selection: The Role of the Basal Ganglia in Language

Jose M. Ceballos

Chair of the Supervisory Committee:

Dr. Chantel S. Prat

The role of the subcortical basal ganglia in language is currently not agreed upon. Research relating these structures to language processes is diverse, and existing theories fail to account for the breadth of findings. To address this, I proposed taking a core basal ganglia neurocomputation, that of reinforcement learning, to explain its involvement in linguistic processes. Lexico-semantics was then used as the level of linguistic processing to test this basal ganglia process. My work consisted of three projects investigating the mechanisms supporting semantic selection: two experiments and one computational simulation using the ACT-R cognitive architecture. Behavioral indices of basal ganglia dual-path learning, or reinforcement learning from positive and negative feedback, showed that semantic selection in the absence of context varies in line with the bias in the estimated value of a specific meaning. The ACT-R computational cognitive model provides a mechanistic explanation of, and causal evidence for, the influence of a reinforcement learning-based action selection system in semantic selection. A self-paced reading sentence task provides further insights into how sentence context and the temporal dynamics of semantic selection process affect basal ganglia-mediated ambiguity resolution. Results from the three projects converge to show that individual differences in learning from both positive and negative feedback through basal ganglia competitive dynamics relate to variability in semantic selection processes. Findings are discussed in light of a general reinforcement learning neurocomputational approach to understanding basal ganglia involvement in language, more generally.

Table of Contents

Chapter 1. Introduction

Research on the neurobiology of language has traditionally focused on contributions of cortical structures to linguistic processes. There are a variety of reasons for this cortex-centric approach witnessed in the field including: (a) an overly simplistic neurobiology of language model rooted in a view of Broca and Wernicke's areas as *the* language regions (Tremblay & Dick, 2016), and (b) the challenges associated with imaging the human subcortex (Forstmann et al., 2017). Together these have contributed to an outsized focus on cortical contributions to human language and cognition.

Converging evidence, however, suggests that the subcortical basal ganglia also contribute to linguistic processes. For example, research on both healthy individuals and those with basal ganglia damage has implicated the basal ganglia in a wide variety of linguistic processes including speech production (Booth, Wood, Lu, Houk, & Bitan, 2007; Fabbro, Clarici, & Bava, 1996; Gil Robles, Gatignol, Capelle, Mitchell, & Duffau, 2005; Senft, Stewart, Bekolay, Eliasmith, & Kröger, 2016; Tettamanti et al., 2005; Wallesch, Henriksen, Kornhuber, & Paulson, 1985), grammatical processing (de Diego Balaguer et al., 2008; Frisch, Kotz, Cramon, & Friederici, 2003; Prat & Just, 2011; Sambin et al., 2012), lexico-semantics (Angwin, Chenery, Copland, Murdoch, & Silburn, 2007; Arnott et al., 2011; Castner et al., 2007; Chenery, Angwin, & Copland, 2008; Copland, 2003; Copland et al., 2012; Copland, Chenery, & Murdoch, 2000; 2001), and even multilingual language control (Friederici, 2006; Hervais-Adelman, Moser-Mercer, & Golestani, 2011; Hervais-Adelman, Moser-Mercer, Michel, & Golestani, 2015; Seo, Stocco, & Prat, 2018; Simmonds, Leech, Iverson, & Wise, 2014; Simmonds, Wise, Dhanjal, & Leech, 2011; Stocco & Prat, 2014; Stocco, Yamasaki, Natalenko, & Prat, 2014). To date, no existing account of the neurobiology of language is able to systematically explain what the role of these subcortical structures is across levels of linguistic processing. Thus, in its current stage, the field suffers from an incomplete understanding of the neural bases of language. In an effort to contribute to a whole-brain model of language functioning, this work focuses on understanding how the basal ganglia nuclei contribute to linguistic processes.

In the following sections I will describe the neuroanatomy necessary for understanding this dissertation work, review evidence for basal ganglia involvement in language, highlight various reasons why our understanding of the basal ganglia in language is so limited, and propose a reinforcement learning neurocomputational framework for understanding subcortical contributions to linguistic processes using semantic selection as a model system.

1.1. Basal Ganglia Neuroanatomy

The basal ganglia are a set of interconnected subcortical nuclei that sit in a privileged position to exert control and modulate the functioning of prefrontal cortex (PFC), as their structural and functional properties give rise to unique signaling dynamics with PFC. The basal ganglia nuclei are characterized by highly intricate inhibitory and excitatory connections. The striatum (composed of the caudate nucleus and putamen) forms is the input structure of the circuit, receiving afferent signals from nearly the entire cortex. This allows the basal ganglia to represent the state of the cortex via its continuous cortical afferent copy. The output structures of the basal ganglia system are the internal segment of the Globus Pallidus (GPi) and the Substantia Nigra pars Reticulata (SNr). Their inhibitory efferents synapse on the ventral nuclei of the thalamus, including the nuclei ventralis anterior (VA) and ventralis lateralis (VL), along with some connections to the midbrain tegmentum (Kemp & Powell, 1971). It is important to keep in mind that the net effect of basal ganglia outputs on the VA/VL nuclei of the thalamus is inhibitory. More specifically, unless inhibited, the GPi engages in continuous release of GABA, leading to a tonic inhibition (60-80 spikes/second in monkeys) of the thalamic nuclei upon which it synapses (Mink, 1996). Thus, the GPi signal modulates PFC through its effect on VA/VL thalamic activity. For a basal ganglia connectivity schematic, see *Figure 1*.

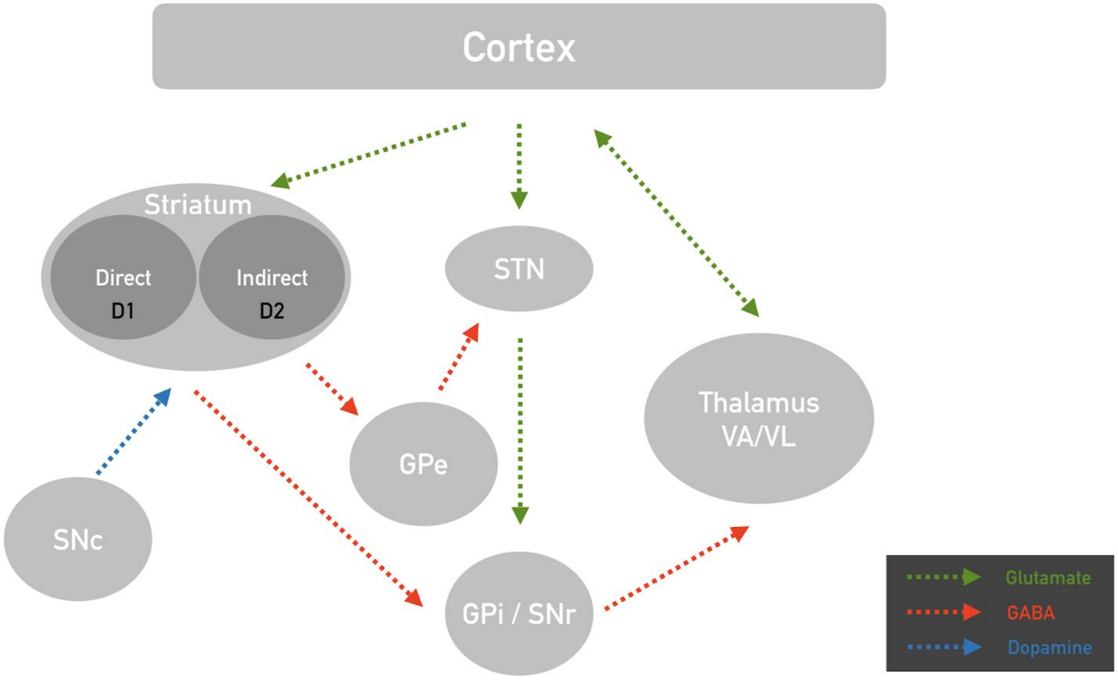

Figure 1. Schematic of basal ganglia connectivity.

This system of inhibitory and excitatory connections between subcortical nuclei is ultimately best understood as forming three distinct pathways, namely: the direct, indirect, and hyperdirect pathways. The direct pathway (also known as the “Go” pathway, see *Figure 2A*) of the basal ganglia makes use of the GPi afferents from the striatum, and the GPi efferents to the VA/VL nuclei of the thalamus. When the striatum receives strong converging afferent “Go” signals matching a specific conjunction or template of activity (Cohen & Frank, 2009; Graybiel, 1995), it fires inhibitory signals directly to the GPi. These inhibitory signals cause the GPi to stop its tonic inhibitory signaling, thus releasing the thalamus from inhibition and increasing thalamic activity. Overall, the net effect of direct pathway activity is to increase signaling to cortical regions (e.g., increase M1 activity, and thus promote motor movement).

The indirect pathway (also known as the “NoGo” pathway, see *Figure 2B*), on the other hand, makes use of several other subcortical nuclei before having its effect on the thalamus. When the striatum engages in “NoGo” signaling, it fires inhibitory signals to the GPe. The now inhibited GPe, which normally engages in tonic inhibition of the STN when active (~70 spikes/second in monkeys), then releases the STN from inhibition (Mink, 1996). STN excitatory signaling synapsing upon the GPi then increases GPi activity, resulting in more VA/VL thalamic nuclei inhibition. Overall, the net effect of indirect pathway activity is to decrease signaling to cortical regions (e.g., decrease M1 activity, and thus stall motor movement). For example, Parkinson's Disease (PD), who have decreased functioning of the direct pathway and an overactive indirect pathway, exhibit pathological blocking of motor cortex which prevents the initiation of a movement.

The hyperdirect pathway (see *Figure 2C*) bypasses the striatum altogether (hence the name “hyperdirect”). Excitatory cortical afferents to the STN cause this nucleus to engage in excitatory signaling to the GPi. This excitatory signaling upon the GPi cause it to increase its inhibitory firing on the thalamic nuclei. The net effect of hyperdirect pathway activity is to decrease cortical signaling.

Furthermore, the activity of the direct and indirect pathways is heavily modulated by the dopamine neurotransmitter system. The substantia nigra pars compacta (SNc) and the ventral tegmental area (VTA) are the main sources of striatal dopamine. Dopamine has opposing effects on two receptor subsystems in the striatum: (1) dopamine D1 receptor binding increases activity in the striatum neuronal subpopulations that make up the direct pathway, and (2) dopamine D2 receptor binding decreases activity in the striatum neuronal subpopulations that make up the indirect pathway. Thus, the net effect of dopamine release by the SNc to the striatum is an increase in thalamic activity to the cortex via inhibition of the GPi. Conversely, in the absence of dopamine, neurons with a high density of D2 receptors (the indirect pathway) become more active, leading to a net decrease in thalamic activity.

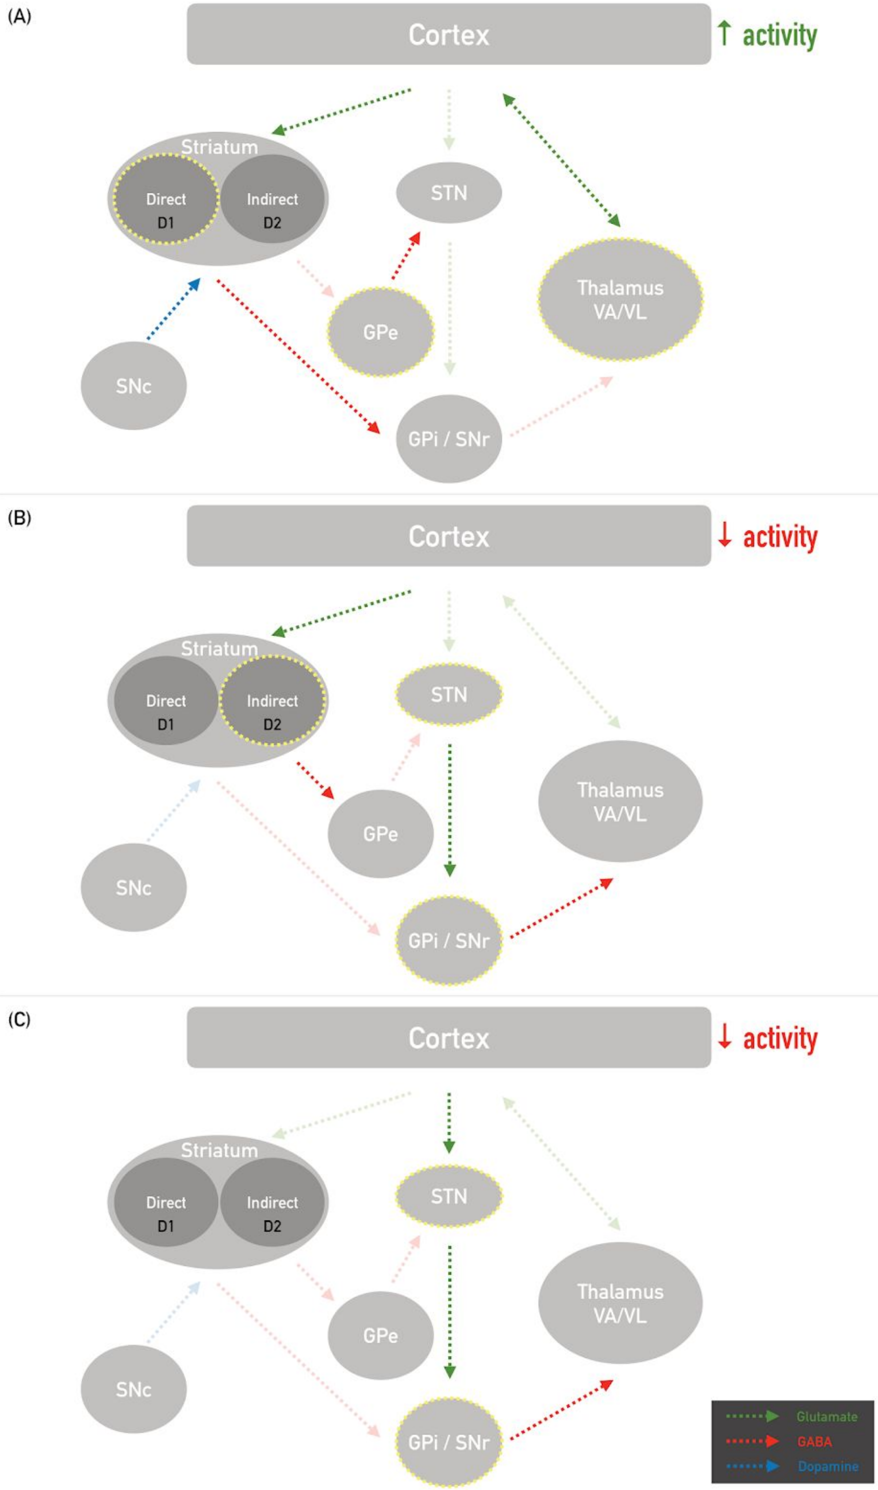

Figure 2. Schematic of the (A) direct, (B) indirect, and (C) hyperdirect pathways.

1.2. Evidence for Basal Ganglia Involvement in Language

Research over the past few decades has provided remarkably diverse evidence that the basal ganglia are involved in language, stemming from neuroimaging research with healthy participants, to behavioral research on clinical populations. To demonstrate this, in a recent review I summarized evidence from 87 research articles describing evidence for basal ganglia contributions to language (Ceballos, Prat, Zeitlin, & Stocco, in preparation). *Supplementary Table 1* includes a list of studies summarized in this review, alongside information regarding whether it was a production or comprehension study, linguistic process studied (e.g., syntax or lexico-semantics), participant population used (e.g., healthy monolinguals or patients with neurodegenerative diseases), research method utilized (e.g., neuroimaging or behavioral), and which basal ganglia nuclei are discussed in light of their contributions to language. In the following sections I will highlight just some of the diverse evidence of basal ganglia involvement in language.

1.2.1. Phonological Processing.

Perhaps the least surprising level of linguistic processing for basal ganglia involvement is speech production and comprehension, as these subcortical structures are well understood in light of their role in motor control. For example, a functional magnetic resonance imaging (fMRI) study showed that the putamen is activated when engaging in naming or word-generation tasks for less-proficient second or third language learned late in life, compared to first or even second language learned early in life (Abutalebi et al., 2013). The authors suggest that native-like fluency reduces the demands on the putamen for speech production, thus suggesting that the basal ganglia play an important role in controlled motor program execution. Furthermore, a speech perception task also elicited putamen activation, as measured by fMRI, when determining if word-pairs rhymed or not (Booth et al., 2007). The authors attribute this basal ganglia involvement to initiation of phonological representations required for phonological segmentation.

Additional evidence comes from patients with genetic mutations that lead to basal ganglia grey and white matter developmental abnormalities. The FOXP2 gene, previously dubbed *the* language gene, is implicated in brain development and is particularly important for grey matter maturation in the striatum (Lai, Gerrelli, Monaco, Fisher, & Copp, 2003; Takahashi, Liu, Hirokawa, & Takahashi, 2003). Research suggests that FOXP2 is crucial for normal synaptic plasticity necessary for early motor-skill learning and early language development (Fisher & Scharff, 2009). Furthermore, research in humans with FOXP2 mutations has also shown that they exhibit abnormally low levels of grey matter in the caudate nuclei. This translates into severe verbal dyspraxia, or a developmental disorder resulting in an inability to coordinate speech processes (Liégeois et al., 2003; Vargha-Khadem, Gadian, Copp, & Mishkin, 2005; Vargha-Khadem et al., 1998; Watkins et al., 2002). Additionally, animal models have provided some comparative evidence, as mice with FOXP2 mutations show decreased dopamine in the striatum, which translates into impairments in vocalizations (Enard, Gehre, Hammerschmidt, & Cell, 2009; Fujita et al., 2008). Similar FOXP2 mutations in songbirds impair vocal learning (Fisher & Scharff, 2009).

1.2.2. Morpho-syntactic Processing.

Children with specific language impairment (SLI) a delay in language acquisition found in children who appear to be normally developing, exhibit grammatical parsing deficits (van der

Lely & Christian, 2000). Research on the neural bases of SLI has suggested that abnormal fronto-striatal systems may be responsible for this acquisition delay (Ullman & Gopnik, 1999). Furthermore, research has demonstrated that basal ganglia damage often results in a pattern of language abnormalities once associated strictly with non-fluent (or Broca's) aphasia (Alexander, Naeser, & Palumbo, 1987). The impairments observed are marked by incorrect use of regular morpho-syntactic rules and are traditionally referred to as "agrammatism" (Ullman, 2004). Additionally, PD patients show abnormal grammatical processing. Specifically, studies using the P600, an electrophysiological component (event-related potential, ERP) known to index sensitivity to grammatical well-formedness (Osterhout & Holcomb, 1992), show reduced sensitivity to grammatical violations, as indexed by smaller or absent P600 components (Friederici, Kotz, Werheid, & Hein, 2003). This abnormal P600 ERP component was also observed in patients with lesions to the basal ganglia caused by subcortical infarct (Frisch et al., 2003; Kotz, Frisch, & Cramon, 2003). Patients with other kinds of basal ganglia deficits, such as patients with Huntington's Disease (HD) with hyperactive direct pathways, show similar impairments in grammatical processing (Lieberman, Kako, Friedman, Tajchman, & Feldman, 1992; Sambin et al., 2012; K. Takahashi et al., 2003).

1.2.3. Lexico-Semantic Processing.

Research has also demonstrated that the basal ganglia are consistently active in lexico-semantic processing (Kuchinke et al., 2011). For example, a study looking at the neural bases of semantic ambiguity processing showed that basal ganglia are reliably activated when reading semantically ambiguous sentences, compared to control sentences with no ambiguity (Mason & Just, 2007). Furthermore, studies have shown that the left caudate nucleus of the basal ganglia becomes significantly active during conditions of high ambiguity (Ketteler, Kastrau, Vohn, & Huber, 2008). Interestingly, the activation of the nuclei is related to the polarity (frequency difference between the dominant and subordinate word meanings) of the ambiguous task words. Also, studies have also shown that the basal ganglia are involved in word generation, but not in nonsense syllable generation (Crosson et al., 2003), suggesting that lexico-semantic retrieval may partly explain basal ganglia involvement during speech production.

1.2.4. Multilingual Language Processing.

Research on second language (L2) acquisition found that caudate nucleus activity predicts L2 literacy skills in children (Tan et al., 2011). These results remain significant after controlling for performance on lexical decision tasks, L2 literacy from the previous year, and measures of intelligence such as nonverbal IQ. Furthermore, a neuroimaging study found that, compared to L1 speech production, the bilateral globus pallidus is more active during L2 speech production (Simmonds et al., 2011), suggesting a critical role in the monitoring of second language phonological mappings. Similarly, a functional imaging study showed that the striatum is involved in the production of novel and L2 sounds, and activation in this region gradually declines as speech becomes more automatic (Simmonds et al., 2014).

The research previously summarized demonstrates the diversity of basal ganglia involvement in language. This diversity observed across production, comprehension, reading, and across multiple levels of linguistic processing have contributed to wildly different hypotheses and theories on the role of these subcortical structures in language. In the following section I will review how, to date, none of the proposed hypotheses provide successful explanations for the evidence discussed previously.

1.3. Existing Theories of Basal Ganglia Involvement in Language

The Declarative/Procedural Theory (Ullman, 2001; 2004; 2016; Ullman et al., 2008) is centered on the role of the declarative and procedural memory systems in linguistic processes. Specifically, Ullman proposes that the declarative system gives rise to the lexicon, while the procedural system is responsible for rule-based linguistic functions such as generative grammar (including syntax and morphology), and the combination of sounds to form natural language (phonology). Importantly, the semantic system is rooted in anterior temporal cortex and the procedural system is born out of PFC-basal ganglia signaling. According to Ullman (2001), the basal ganglia "procedural [system] subserves the implicit learning and use of a symbol-manipulating grammar across subdomains that include syntax, morphology and possibly phonology" (pg. 718). As such, the basal ganglia forms the basis of the combinatorial property of language and is responsible for mapping (during learning and comprehension) and establishing (during production) order between phrases, words, and morphemes.

In grounding his theory on the declarative and procedural memory systems of the brain, Ullman generates an account of the neurobiology of language that dichotomizes the neural bases of semantics and grammar. This account allows itself to account for the linguistic processes involving the 'procedural functions of the basal ganglia' such as speech production, grammar, and cognitive processes dependent on sequencing operations. Naturally, such a dichotomization of processes can lead to challenges, such as when there is mounting evidence that this "procedural system" is also heavily involved in linguistic processes otherwise associated with declarative knowledge. For example, and crucial to theories on the neurobiology of semantics, Ullman's model leaves questions unanswered about selection processes during language comprehension. The basal ganglia have been routinely documented as being involved in semantic ambiguity, where word meaning alone is insufficient to properly select an appropriate meaning. According to the declarative/procedural model, the declarative system is responsible for the learning, representation, and use of all lexical information. Thus, the strict involvement of basal ganglia in procedural processes is at odds with semantic processing under instances of ambiguity resolution.

A contrasting example to Ullman's theory involves a theory on the neural basis of language control (Friederici, 2006). Specifically, this theory extensively implicates the basal ganglia, and in particular the left caudate nucleus, in controlled language processes. Friederici's theory is centered around the distinction between controlled linguistic processes, or L2 learning and highly automatic processes, such as L1 reading of high-frequency words, proposing that controlled processes, in particular, rely on the basal ganglia. This control theory suggests that the left caudate nucleus may be generally involved in the adaptation to varying linguistic processing demands. Specifically, rather than assuming the caudate is involved in switching from one word meaning or one language to another (e.g., Crinion et al., 2006), Friederici summarizes a set of results that support the notion that the left caudate is recruited under circumstances where language processing is less automatic and requires greater control. This is consistent with subsequent research by Prat and colleagues, showing that in a sentence reading task, the degree to which basal ganglia activation increased with greater task demands was correlated with executive abilities (Prat & Just, 2011). Some additional evidence reported to support this view is the fact that monolingual readers activate the left caudate more when reading words learned late in life, compared to those learned early in life (Bradley, King, & Hernandez, 2013). Given that speakers have more practice with words learned early in life (i.e., more automatic), this may be reflective of the left caudate being recruited for control-related linguistic functions.

Friederici's theory of the role of left caudate during the recruitment of top-down control for language use can be seen as providing an orthogonal view to that of the declarative/procedural model. Rather than accounting for a breadth of linguistic phenomena (e.g., speech production, syntax, morphology, etc), the theory for controlled language processing aimed to explain language control in the bilingual (and monolingual) brain. Indeed, the mechanisms proposed by Friederici can be extended to account for lexical selection via the same processes that control language use. Evidence suggests that the basal ganglia are used under situations of lexical ambiguity resolution (Copland, 2003; Copland et al., 2000; 2001; 2012; Mason & Just, 2007), where greater control is necessary to resolve competition between lexical representations. However, this control theory is not all-encompassing; a variety of phenomena described herein fall outside of the scope of "controlled processing", yet still reliably recruit the basal ganglia. For example, I have summarized evidence that many automatic rule-based behaviors covered under Ullman's model (e.g., articulation and canonical grammatical parsing) rely upon basal ganglia contributions. Such processes go unaccounted for in a theory of "top-down" or controlled processing. In fact, although they are not completely orthogonal, the phenomena proposed by the two views previously presented tend to be largely non-overlapping.

Kotz and colleagues (Kotz & Schmidt-Kassow, 2015; Kotz & Schwartze, 2010; Kotz, Schwartze, & Schmidt-Kassow, 2009) posit yet a third role for the basal ganglia in language in a theory of Sensory Predictability. In their work, the authors claim that the underlying process that ties together the variety of non-motor phenomena associated with basal ganglia functioning (e.g., syntax, attention, beat perception, neural timing) is its role as a timekeeper; "the basal ganglia in concert with cortical structures function as a 'pacemaker' that supports synchronization of internal and external cues" (pg. 986, Kotz et al., 2009). This pacemaker is critically involved in predicting when events will occur in a sequence. According to Kotz and colleagues, the basal ganglia work jointly with the pre-SMA to support the domain-general processes of sensory predictability in speech comprehension. With respect to syntax, the cortico-subcortical network is responsible for establishing temporal coherence during syntactic processing. This temporal coherence then gives rise to the anticipation and prediction of upcoming information during speech comprehension. Thus, it is the role of the pre-SMA-basal ganglia system as a pacemaker, which establishes temporal cohesion between internal and external states, that allows for the successful sequencing that is the hallmark of syntactic processing. The authors argue that dysfunctions in syntactic processing witnessed in patients with basal ganglia damage are largely epiphenomenal. More specifically, damage to the pre-SMA-basal ganglia circuit, resulting in disruptions in the ability for the brain to synchronize and re-adjust internal oscillations with external stimuli, prevents the coupling of internal grammatical predictions with reliable external speech cues.

While the Sensory Predictability theory may neatly account for many of the linguistic phenomena that involve sequencing (e.g., speech production and syntactic processing), it is less clear how it may account for other phenomena. For example, it is unclear how this theory would account for findings involving abnormalities in lexical priming observed using single written word presentation tasks (e.g., Copland et al., 2001; 2003). In this Sensory Predictability theory, the signal used to generate internal predictions comes directly from the acoustic information found in incoming speech waves. In its current stage, the theory is unable to account for evidence found across various language domains, including written or even sign language (Robertson, Corina, Ackerman, Guillemin, & Braun, 2004).

1.4. Current Knowledge Gaps

Despite a growing body of literature suggesting that the basal ganglia play a crucial role in language processing, to date, there is no comprehensive theory that fully explains basal ganglia involvement in linguistic processing. There are at least two factors that contribute to this gap in the literature. Because of its seemingly privileged role in human communication, research on the neurobiology of language suffers from a lack of comparative research. So despite the fact that the basal ganglia are an evolutionarily ancient set of neural structures present in virtually all vertebrate systems (Reiner, 2009), and animal models have proven useful at expanding our understanding of the neurobiology of communication (Enard et al., 2009; Fujita et al., 2008), many language neurobiologists have overlooked these subcortical structures with respect to their role in human language.

The second reason involves the way in which the basal ganglia have traditionally been studied by the relatively few language neurobiologists who have focused on these subcortical structures. As said best by Dr. Miriam Bassok during one of her Core Concepts in Cognitive Science lectures at the University of Washington, much of modern-day cognitive neuroscience has an obsession with “blobology.” This is particularly true of language neurobiology, which has an extensive history of placing people into a Magnetic Resonance Imaging (MRI) machine, administering a language task, seeing which areas “light up,” and being quick to ascribe a specific cognitive function to whichever blobs remain after statistical correction. While this example may be seen as somewhat hyperbolic or an overgeneralization of functional imaging, this phenomenon has been previously discussed in neuroimaging research as the “imager’s fallacy” (Henson, 2005). This approach has found its way into investigations of basal ganglia contributions to language functioning and it remains, in my opinion, one of the biggest barriers to advancing our understanding of the role of these structures in language. For example, the level of description of many existing theories ranges from “the basal ganglia is [sic] crucial for meaning comprehension” (Ali et al., 2009), to “the caudate is involved in abstract linguistic processes, while the putamen is involved in motor processing of linguistic programs” (Gil Robles et al., 2004). At this level of description, these theories fail to provide unifying explanations for the myriad of findings previously described herein. In the following quote, Yin & Knowlton (2006) highlight why it is important to consider the basal ganglia as a functional network with respect to its role in human behavior.

> Although each of these [basal ganglia] components, by virtue of characteristic physiological properties, has unique ‘computational’ properties, at the behavioral level it is the integrated functioning of a distributed network comprising various components that is important. This point is worth emphasizing, as systems neuroscience is often dominated by attempts to localize psychological functions without regard for the actual functional circuitry of the brain. Not only do the psychological functions lack operational specificity, but the anatomical entities that are said to subserve such functions also lack the requisite circuitry. (Yin & Knowlton, 2006)

Thus, the current dissertation aims to address a critical gap in the field of the neurobiology of language by making use of well-understood basal ganglia computations to account for its role in linguistic ability.

What evidence do we have to suggest this approach may be a promising one? Strong evidence comes from investigations of working memory capacity (WMC), which have capitalized on the functional network signaling dynamics between the basal ganglia and PFC to create robust explanations of this cognitive construct. Specifically, research on the role of the basal ganglia in WMC suggests that the direct, indirect, and hyperdirect pathways are crucially involved in the maintenance, updating, and filtering of representations in PFC. Specific mechanisms for such processes have been proposed by a variety of experimental and theoretical explorations into the neural bases of WMC. While some of these investigations focus on the active signaling dynamics necessary for maintenance of information in working memory (e.g., Frank et al., 2001; Hazy et al., 2007), and others focus on the selection mechanisms that give rise to capacity WMC constraints (e.g., McNab & Kingberg, 2007), all PFC-basal ganglia investigations of WMC are ultimately rooted on the neurocomputational principles of competitive signaling in the basal ganglia pathways. Thus, it is reasonable to hypothesize this neurocomputational framework, which makes use of our understanding of the basal ganglia as a network that dynamically influences PFC, will be a fruitful approach for this dissertation project.

1.5. A Fundamental Basal Ganglia Neurocomputation

At first glance, it may be tempting to try to explain basal ganglia involvement in language using the same *exact* kind of neurocomputation used by other investigations (e.g., those employed by WMC models). This may prove to be problematic because, while there is a general agreement that the basal ganglia influence PFC functioning, computational models and theories vary dramatically in their implementation of how this occurs. Take for example two canonical computational models of PFC-basal ganglia functioning that provide neurobiologically plausible accounts of WMC; the PFC-Basal Ganglia Working Memory model (PBWM; Frank, Loughry, & O'Reilly, 2001; Hazy, Frank, & O'Reilly, 2007), and the Conditional Routing model (CR; Stocco, Lebiere, & Anderson, 2010). Both models adhere to the widely accepted view that the basal ganglia function in a brake-accelerator fashion by selecting for (accelerate) and against (brake) specific cortical signals (Graybiel, 2000). In the PBWM model, the robust maintenance of information associated with WMC is born out of the direct and indirect pathways functioning as "gates" that write into PFC "slots." Direct pathway "Go" signaling is responsible for allowing information into these PFC slots. Indirect pathway "NoGo" signaling, on the other hand, prevents representations from being overwritten in PFC. Thus, these two pathways allow for information to be selectively encoded and protected in PFC, giving rise to WMC operations. In the CR model, basal ganglia pathways play a role in routing information from a source to a destination and enable flexible behavior and cognition. Direct pathway "Go" signaling enables the transfer of a signal to all of the destinations associated with the source of a specific representation. Indirect pathway "NoGo" signaling then narrows down destination regions to all but one. In this way, the CR model posits that WMC robust maintenance is born out of PFC representations where the source and destination signaling is the same. So while a network-level neurocomputational approach may prove useful in accounting for basal ganglia involvement, the problem space for possible specific implementations remains intractably large.

A careful look at these computational models highlights the fact that, across varying implementations and mechanisms, the basal ganglia make use of feedback signals that enable learning of when and how to use direct and indirect pathway signaling to affect PFC. In other words, the "common denominator" neurocomputation observed across models and theories of PFC-basal ganglia is that of *reward processing*. This is worth emphasizing because, in order for break-accelerator dynamics to successfully modulate cortical signaling in a way that is behaviorally advantageous, the basal ganglia need to be privy to information on the outcomes of such signaling. Thus, at its most fundamental level, it is a focus on the computational mechanisms that give rise to *which* cortical signals to modulate, rather than *how* they are modulated, that may provide a successful account of basal ganglia involvement in language. Fortunately, there is an abundance of research implicating these subcortical nuclei in reward-mediated learning.

1.5.1. Basal Ganglia Reinforcement Learning.

Perhaps the most general and uncontroversial framework used to describe the function of the basal ganglia is that of reinforcement learning (Sutton & Barto, 1998). In its simplest form, a reinforcement learning agent aims to operate inside of an environment in a way that its actions result in the highest possible predicted reward. Thus, a reinforcement learning agent is continually seeking to maximize a cumulative reward from its environment through careful mappings between environmental state and action pairs. The rewarding values of stimuli and actions are learned through the comparison of successive estimates of future rewards, and this comparison yields a scalar term, or Reward Prediction Error (RPE). The RPE is then used to adjust the internal values of the reinforcement learning agent. Specifically, the updated value Q' of an action a_t in response to a state s_t at time *t* is calculated by comparing the previous estimates of Q at time *t* and t+1:

$$Q'(s_t, a_t) = Q_t(s_t, a_t) + \alpha[R_{t+1} + \gamma Q(s_{t+1}, a_{t+1}) - Q(s_t, a_t)]$$

Here, R_{t+1} is the reward (if any) that follows an action at time t+1, $0 < \gamma < 1$ is a temporal discount factor, and a_{t+1} and s_{t+1} represents the next stimulus and action. The quantity $[R_t + \gamma Q(s_{t+1}, a_{t+1}) - Q(s_t, a_t)]$ is the RPE, which captures the difference in reward predictions before and after the action a_t has been performed. The degree to which the RPE affects learning is modulated by the learning rate parameter α, a free parameter that generates variability in models. Higher values of α are related to faster learning, and this parameter has been previously used to capture individual differences in task learning. Furthermore, it is important to emphasize that Q-values are estimates of future rewards based on the frequency and magnitude of rewards. Because actions are chosen on the basis of Q-values, behavior will ultimately reflect internal estimates of frequencies of rewards. Crucial for this dissertation, these estimates may sometimes be systematically biased depending on the mechanisms giving rise to reinforcement learning.

Reinforcement learning accounts have been successful in providing the necessary framework for studying the neurocognitive bases of reward-mediated learning. Single-cell recording studies in rodents and primates have provided ample evidence that reinforcement learning equations provide a surprisingly effective description of cell activity in the basal ganglia. In particular, the firing of midbrain dopamine neurons matches the RPE signal during learning (Schultz, Dayan, & Montague, 1997). This finding spurred the development of computational models using the basal ganglia as the neurobiological bases of reward-mediated learning, with reinforcement learning as the most fundamental neurocomputation mediating action selection and PFC biasing.

1.5.2. Basal Ganglia Reinforcement Learning and Language.

The three theories discussed previously (see section *1.3. Existing Theories of Basal Ganglia in Language*) illustrate the diversity of theories of the role of the basal ganglia in language. Additionally, these three theories highlight how challenging it has been to find a description of the basal ganglia in language that may successfully generalize across the seemingly unrelated pieces of evidence that exist to date. As previously mentioned, this issue likely stems from existing efforts focusing on (a) the contributions of specific nuclei to narrow functions (e.g., "the caudate does syntax"), and (b) the multitude of ways in which basal ganglia action selection can be implemented in computational models. Thus, a network-level approach making use of a core (or common) neurocomputation will provide an integrative view across the largely non-overlapping theories presented beforehand.

The first two theories, Ullman's Declarative/Procedural and Friederici's Control theory can together be understood under the traditional accelerator-brake (Graybiel, 2000), or action selection view of the basal ganglia, which is born out of reinforcement learning neurocomputations (see section *1.3. A Fundamental Basal Ganglia Neurocomputation*). The Declarative/Procedural theory argues for the involvement of basal ganglia in the processes involving the combinatorial properties of language, such as generative grammar. In order for combinatorial systems to properly generate the correct outputs (e.g., "jump" + [past tense marker] = "jumped"), selection between competing grammatical markers and rules needs to occur. This function can clearly be ascribed to basal ganglia neurocomputations, and in fact, previous research has used a basal ganglia reinforcement learning-based action selection model to give rise to generative grammar (Kriete, Noelle, Cohen, & O'Reilly, 2013). The Control theory can also fundamentally be broken down using this lens. When looking at what instances makeup moments of “high control” for basal ganglia recruitment (e.g., using the contextually appropriate language or word), it is apparent that the core issue is one of selecting the correct cortical representation between competing ones. Furthermore, the theory of Sensory Predictability *is* essentially a theory of learning and prediction born out of basal ganglia reinforcement learning (e.g., Schultz et al., 1997). While this theory has focused on phonetics and morphosyntax as the levels of linguistic processing to study and explain, the core processes of live prediction during language processing should generalize more broadly.

It is worth highlighting that across all of these theories, one of the core requirements is that the basal ganglia is operating in a way that is advantageous to cognition and behavior by selecting or biasing the *correct* information in PFC. This highlights the importance of taking a reinforcement learning perspective that, while remaining agnostic to the specific instantiations of action selection (e.g., writing into PFC “slots” versus routing of cortical representations), makes predictions about how individual differences in *learning* and *prediction* will be critically related to variability in linguistic ability. In the following section I will review recent advances in the field’s understanding of basal ganglia-mediated learning that have proven critical in accounting for nuanced individual differences in human cognition.

1.5.3. Dual-Path Reinforcement Learning.

To date, the majority of simulations making use of basal ganglia dynamics in reinforcement learning are modeled using a Temporal Difference (TD) method. TD is a reinforcement learning algorithm which learns without explicit knowledge about the dynamics observed in its environment (it is, in reinforcement learning jargon, “model-free”) by updating subsequent estimates of reward prior to knowing a final outcome (Sutton & Barto, 1987). This is problematic for two key reasons: (a) A TD method overlooks the independent contributions of the direct and indirect pathways to reward-mediated learning, and (b) TD models fail to capture important individual differences in human behavior that are known to arise from pathology, neurobiological composition, and pharmacological intervention that may affect either basal ganglia pathway (Cohen & Frank, 2009). In an effort to address this, a recent computational modeling approach was proposed. Specifically, this approach capitalizes on the opposing pathway dynamics of the basal ganglia and implements a dual-path architecture for reward-mediated learning in model systems that have traditionally used TD learning (Stocco, 2018). Importantly, this dual-path model provides evidence for basal ganglia-mediated learning that arises from individual differences in sensitivity to positive and negative feedback. In other words, traditional TD learning methods that only learn from positive feedback do not consider critical information contained in negative feedback, something the basal ganglia do naturally through their competing pathway dynamics. Furthermore, this model showed how sensitivity to both types of feedback makes for a robust reinforcement learning agent that better accounts for human performance in canonical tasks of implicit learning (Stocco, 2018). This model provides important advancements in understanding how the basal ganglia pathways contribute to learning processes, and in turn, provides a robust framework for exploring core basal ganglia neurocomputations and their contributions to language. Next, I will discuss the use of semantic selection as the language process of choice for this dissertation work and elaborate why a dual-path reinforcement learning account may prove insightful in our understanding of semantics, and ultimately the neurobiology of language.

1.6. Semantics as a Model System

The current dissertation focuses on the role of basal ganglia reinforcement learning mechanisms in lexico-semantics, generally, and semantic selection processes more specifically. I will refer to these processes as "semantics" throughout the dissertation but acknowledge that the term semantics is much broader and spans levels of comprehension and multiple construction processes. The choice to focus on lexico-semantics was both pragmatic (as it provided a manageable "chunk" of linguistic phenomena to address in a dissertation) and theoretically driven. From a theoretical standpoint, semantic provides an interesting test case because it is a well understood and operationalized case of selection processes that make use of selection processes likely to originate from the basal ganglia pathways. In the following sections, I will discuss the neurocognitive bases of semantic selection processes necessary for ambiguity resolution.

1.6.1. Semantic Ambiguity Resolution.

Ambiguity is nearly ubiquitous in language, as most words can refer to multiple concepts, or at least subtly different senses of related concepts. For example, the word "hot" can refer either to temperature or to food spiciness. While most commonly referred to as lexical ambiguity, the term semantic ambiguity will be used herein to highlight the fact that the phenomenon of interest is ambiguity in the meaning of a word, and its neural representation, rather than its written or grammatical form. Language use in a naturalistic setting rarely results in semantic ambiguity, as the conversational context (e.g., topic of conversation, role and status of interlocutor, location, nonverbal communication cues) along with word frequency (the relative frequency with which speakers encounter a word in a language) provide enough information to disambiguate between various possible word meanings. However, cases of semantic ambiguity do arise in conversational settings and are often encountered in written form such as news headlines, puns, poetry, and novels. The ability to properly disambiguate an input into the contextually appropriate represented meaning is key for listening and reading comprehension. More importantly, this process provides details on fundamental neurocognitive mechanisms, such as the contextual integration of information, inhibition, and selection processes, used to manage simultaneous conflicting neural representations that are at odds with the task goal of accurate transfer of information in communicative settings.

Semantic ambiguity can arise in a variety of different ways. The first class of ambiguity arises from words that have different unrelated meanings. For example, "bark" can refer to the sound a dog makes or the outermost layer of a tree. In this case, both meanings of "bark" constitute a true homonym but are also homographs and homophones (same spelling, and same sound, respectively). Furthermore, words can be encountered in contexts where only the written form is ambiguous (e.g., the homographs for "lead"), or only the spoken form is ambiguous (e.g., the homophones for "be/bee" or "seam/seem"). The second class of ambiguity arises from words where semantic forms are related in meaning. For example, in the sentences "the child fell while running," "the lovers fell in love," and "the man fell for the lie," the word "fell" refers to different concepts that partly overlap in meaning. This type of semantic ambiguity, called polysemy, is more common in natural or everyday language use.

1.6.2. Cognitive Processes Shaping Ambiguity Resolution.

The cognitive mechanisms supporting semantic ambiguity resolution are best understood by exploring psycholinguistic research and theories on the dynamics of semantic information and its representation. The Reordered Access Model (RAM; Duffy, Morris, & Rayner, 1988; Duffy, Kambe, & Rayner, 2001) states that when a listener or reader first encounters a word with multiple meanings, all meanings are activated and available in parallel (Duffy, Morris, & Rayner, 1988). Evidence for this comes from cross-modal semantic priming studies in which the probe was an ambiguous homograph. This research found that responses to target words related to the ambiguous homograph were faster than responses to unrelated words (Seidenberg, Tanenhaus, Leiman, & Bienkowski, 1982). Furthermore, if encountered in isolation or in a highly ambiguous context, an ambiguous word will be automatically disambiguated towards the highest frequency meaning within hundreds of milliseconds, reflecting processes of semantic selection (Swinney, 1979). However, if an ambiguous word is encountered following a strong biasing context towards the dominant or more frequent meaning, only this contextually-relevant dominant word meaning will be available (Tabossi & Zardon, 1993). Thus, when ambiguous words are encountered, all meanings are initially activated, but this activation is modulated and constrained by factors such as meaning frequency and sentence context.

1.6.3. Neural Processes Shaping Semantic Ambiguity Resolution.

Research has converged on the view that parallel activation of multiple meanings is ultimately narrowed down by the left inferior frontal gyrus (LIFG), giving rise to the notion that the LIFG engages in processes of semantic selection. This view has been born out of functional imaging studies which have shown that the LIFG, and specifically the pars triangularis, pars opercularis, and pars orbitalis show greater activation for ambiguous sentences, relative to a control and unambiguous baseline (Mason & Just, 2007; Rodd, Davis, & Johnsrude, 2005; Rodd, Johnsrude, & Davis, 2010). In fact, the LIFG has been the region most consistently shown to have a strong blood-oxygen-level-dependent (BOLD) response for semantically ambiguous sentences, reported across nearly all studies of semantic ambiguity resolution (Vitello & Rodd, 2015). Furthermore, this effect is known to be highly reliable, as over 80% of individuals show ambiguity-related changes to the BOLD response specifically in the pars triangularis of the LIFG during both sentence and single-word ambiguity tasks (Vitello, Warren, Devlin, & Rodd, 2014).

The Conflict Resolution Account (CRA; Schnur et al., 2009) provides the most widely-accepted explanation for why the LIFG is involved in semantic ambiguity resolution. Specifically, the CRA posits that the LIFG resolves competition that is born out of multiple active cortical representations. When there is a dominant but irrelevant response, or when there are multiple equally probable responses to be selected, the LIFG makes use of additional top-down modulatory signals to narrow down, and ultimately select, the appropriate response. This raises the critical question on the origin of top-down modulatory signaling to LIFG. To date, research on the neurobiology of semantic ambiguity resolution has given very little focus to understanding or characterizing this biasing signal, and existing work largely adheres to the Bias Competition Framework (Miller & Cohen, 2001) which posits that this bias originates from PFC (Vitello & Rodd, 2015). However, critical to this investigation is research summarized previously that demonstrates that the PFC itself is modulated by the basal ganglia (Cohen & Frank, 2009; Frank et al., 2001; Hazy et al., 2007; Ketteler et al., 2008). Furthermore, research on the functional and anatomical properties of the PFC-basal ganglia network has shown that two of the five main cortico-striatal-thalamocortical loops project directly to lateral prefrontal regions, including dorsolateral PFC and lateral orbitofrontal cortex (Alexander & Crutcher, 1990). Thus, the basal ganglia possess the functional, anatomical, and computational properties necessary to provide biasing signals to LIFG during semantic ambiguity resolution. Based on this, I make the novel prediction that individual differences in basal ganglia reinforcement learning, which give rise to action selection dynamics, will relate to variability in performance in tasks measuring semantic ambiguity resolution.

What evidence exists that may suggest that basal ganglia reinforcement learning may play a key role in semantic ambiguity resolution? Research on patients has shown that priming effects are significantly longer lasting when there is damage to either of the two basal ganglia pathways. Specifically, PD patients have access to multiple word meanings associated with an ambiguous homographs (Copland, 2003; Copland et al., 2000; 2001). Access to multiple meanings, including low frequency (or subordinate) meanings, remains even after long delays that result in semantic selection for healthy controls. Similarly, Huntington's Disease (HD) patients, who exhibit the opposite type of functional pathway damage of PD patients and instead have decreased functioning of the indirect pathway and an overactive direct pathway, also exhibit prolonged semantic priming and abnormalities in semantic selection (Chenery, Copland, & Murdoch, 2002). While these findings have been previously characterized as the basal ganglia *directly* engaging in inhibition (for PD) or selection (for HD) of competing semantic cortical representations, this pattern of results across opposing action selection pathways can be viewed as the result of single-path or TD-like learning and prediction in the brain. Like TD models of probabilistic learning, patients with only one functioning basal ganglia pathway produce inaccurate estimates of low probability stimuli (Van Hasselt, Guez, & Silver, 2016). Importantly for this investigation, TD reinforcement learning methods show large overestimations of the values of low-reward options, even when you reverse TD learning (i.e., only negative rewards). This bias is only fixed when two-path systems are implemented (Van Hasselt, Guez, & Silver, 2016). Thus, this would translate to both PD and HD patients have a biasing signal to LIFG that is stronger for the subordinate meaning, relative to healthy controls. Importantly, the basal ganglia biasing signal during semantic selection can be interpreted as reflecting linguistic statistical knowledge, in this case, that of meaning frequency. This account provides a plausible and testable account for how the basal ganglia are involved in semantic selection processes for ambiguity resolution.

1.7. Aims

A primary goal of this dissertation is to expand the understanding of the subcortical neurobiological processes that support human linguistic ability. This study will elucidate the contributions of the basal ganglia structures to human semantic processes, and specifically to semantic ambiguities and the neurocognitive mechanisms for resolving them. One of the reasons for focusing on semantic ambiguity resolution stems from this being consistently identified as a process that is dependent on basal ganglia functioning but lacking well-specified models to account for the observed phenomenon. A secondary goal of this study is to understand the nature of basal ganglia neurocomputations involved in language processes and relate them to non-linguistic processes. Specifically, this secondary goal aims to test the hypothesis that basal ganglia contributions to human cognition are qualitatively similar across linguistic and nonlinguistic functions. Both of these aims will be tested by using a well-specified theory of basal ganglia reinforcement learning that has been previously used to explain its role in non-linguistic cognitive processes.

1.8. Overview

To test the novel hypothesis that individual differences in basal ganglia dual-path reinforcement learning will relate to variability in semantic ambiguity resolution I employ the following three-step approach. First, I use a semantic relatedness word task to measure how often people select a particular meaning of an ambiguous homograph, which should vary in line with their reinforcement learning estimate of utility of selecting a specific word-meaning. Second, I use a computational cognitive model of basal ganglia dual-path learning and action selection to study the reinforcement learning mechanisms associated with semantic selection processes during ambiguity resolution. Finally, I use a more naturalistic sentence task to better understand the effects of sentence context and semantic selection dynamics in ambiguity resolution by measuring self-paced reading speed.

Chapter 2. Experiment One: Context-free Semantic Processing

2.1. Operationalization of Context-free Semantic Processing

As discussed previously, context plays an important role in semantic processing. This is particularly apparent in the constraining effects context plays during semantic ambiguity resolution. Although such constraint processing is prevalent in natural language processing, it increases the problem space for modeling and understanding individual differences in semantic ambiguity resolution. Thus, the first experiment explores these processes in a context-free paradigm.

Biased homographs, or ambiguous words that have relatively higher (dominant) and lower (subordinate) frequency meaning associates, were chosen as the ambiguous critical words for this investigation. This manipulation follows a similar experimental paradigm used by Copland and colleagues (Copland, 2003). The use of biased homographs (as opposed to balanced homographs, where both meanings occur with relative equal frequencies) facilitates the exploration of automatic ambiguity resolution processes for various reasons. For example, the strength of activation of dominant and subordinate meaning associates varies following the onset of the ambiguous homograph, a process that has previously shown to be abnormal in patients with subcortical damage or neurodegeneration. Previously, a lexical priming investigation found large individuals differences in the neurocognitive dynamics that shape lexical priming in basal ganglia clinical populations, relative to healthy controls (Copland et al., 2001). Specifically, healthy participants showed no traces of subordinate word activation following a long delay between a homograph and semantic target, thus reflect automatic semantic ambiguity resolution towards the dominant or highest frequency meaning. Parkinson's Disease (PD) patients on the other hand, who have decreased dopaminergic signaling resulting in hyperactivity of the basal ganglia indirect pathway, exhibit an indeterminate or longer-term activation of the multiple competing representations. Thus, this study motivates the use of biased homographs in a paradigm that may provide further insight into the effects of the basal ganglia pathways in lexically ambiguous priming.

A varied interstimulus interval (ISI) between the critical word and the semantic target was used in order to better understand individual differences in the time course of memory activation of the dominant and subordinate meanings in healthy populations. While Copland and colleagues made use of a short ISI of 200 ms and a long ISI of 1250 ms, at the long ISI they found a homogenous performance by the healthy control group, with healthy controls showing no signs of access to the subordinate word meaning. When relating these findings to other investigations of access to subordinate word meanings following varied ISIs (Gernsbacher, Varner, & Faust, 1990), a shorter duration of the long ISI condition can indeed give rise to vast individual differences in task performance, even in a healthy population. In other words, very long ISIs ensure that essentially all individuals select for a specific meaning, while shorter times may capture important variability in this semantic selection process. For this reason, the current experiment made use of a short ISI of 850 ms.

A learning and prediction account of the basal ganglia in semantic ambiguity predicts that there is a crucial role played by both the direct and indirect pathways in learning about statistical regularities in the environment, and thus ultimately in accurately evaluating the predictive utility of linguistic information. Therefore, relatively equal direct and indirect pathway contributions to learning and prediction will result in a behavioral pattern of results that is qualitatively different from that of only high direct or high indirect pathway functioning. In the latter case, this type of learning will resemble the learning from traditional single-path (e.g., TD algorithms) reinforcement learning agents showing a misrepresentation of low-frequency probabilities.

2.2. Methods

2.2.1. Participants.

Informed consent was obtained from participants prior to the experiment, as outlined by the University of Washington Institutional Review Board. Participants were recruited using the Psychology Department's Participant Research Pool (PRP) and all participants were compensated with course credit for their participation. Data were collected from 140 healthy participants (66 females, mean age = 19.4 years). All participants were monolingual, defined as speaking English fluently and having no L2 fluency. Screening for admission was done directly through the PRP, and only monolingual participants had the opportunity to sign up for the study. Seven participants were removed from analysis due to issues with data collection or poor task performance. All participants completed the experimental tasks in four pseudo-randomized orders to control for possible order effects induced by the WPT and PSS task length.

2.2.2. Context-free Semantic Judgment Task.

The Word-Pair Task (WPT) was designed to measure the availability of dominant and subordinate word meanings following the presentation of homographs with multiple meanings. The homographs used shared both phonetic and orthographic forms across both word meanings, making them "true homographs" (e.g., "Bat"). The homograph and semantic target words were presented in the center of the screen, one at a time, separated by an ISI of either 150 ms (short) or 850 ms (long). Prior to starting the task, participants were asked to place their right index finger on the "P" key of the keyboard, and their left index finger on the "Q" key of the keyboard. Participants were then asked to respond with a button press if the target word was related or unrelated to the critical word. Key mappings for related and unrelated responses were counterbalanced.

There were two conditions of interest, (a) homograph & dominant meaning target, (b) homograph & subordinate meaning target, and two control conditions (c) non-homograph & related word target, and (d) non-homograph & unrelated word target. These four conditions will be referred to as Dominant, Subordinate, Related, and Unrelated (respectively) from here on for simplicity purposes. Participants completed 100 total word-pair trials, where 20 belonged to the Dominant condition, 20 to the Subordinate condition, 30 to the Related condition, and 30 to the Unrelated condition. Word frequency meanings were obtained from (Twilley, Dixon, Taylor, & Clark, 1994), and subordinate words were defined as having a relatedness frequency of less than 0.3 but greater than 0.1 on a 0-1 scale. The minimum cutoff of 0.1 was implemented in order to avoid using very obscure words. Dominant words had a relatedness frequency of greater than 0.7. Since each homograph is associated with two meanings, but each homograph was presented once for each participant, two WPT versions were created. In one version, the dominant meaning of a homograph was used (e.g., version A contained “Bank”" & “Money”) while the other version used the subordinate meaning (e.g., version B contained “Bank” & “River”). The two lists were counterbalanced for word frequency, word length, and syllable length.

2.2.3. Probabilistic Stimulus Selection Task.

The Probabilistic Stimulus Selection (PSS) task is an iterative, two-alternative, forced-choice, decision-making paradigm first introduced by (Frank, Moustafa, Haughey, Curran, & Hutchison, 2007). In this task, participants are repeatedly asked to select one of two stimuli presented on the screen. Participants are also told that some of their choices would result in “success”, and some of them would result in “failure”, depending on which stimulus they choose. Feedback on the outcome of their decision is presented immediately after participants select a stimulus. To encourage participants to avoid explicit strategies (such as rote memorization of each stimulus’ history of successes), stimuli are implemented as complex shapes that are difficult to verbalize—typically, Hiragana characters presented to non-Japanese speaking participants. Unbeknownst to participants, each stimulus has a predefined "success" probability. Six stimuli in total are used in the experiment, with success probabilities of 80, 70, 60, 40, 30, and 20 percent. In the first phase, the stimuli are divided into fixed pairs (80 and 20, 70 and 30, 40 and 60), with the highest probability stimulus always paired with the lowest probability one, then second higher stimulus paired with the second lowest one, and the third highest probability stimulus paired with the third lowest one.

Two values are calculated from the test phase of the PSS task: *Choose* accuracy, which represents the accuracy in choosing the most rewarding stimulus over others; and *Avoid* accuracy, that is, the proportion of times in which participants avoid the least rewarding stimulus. The six stimuli are indicated with the letters A, B, C, D, E, and F, with A being the most rewarding stimulus and B the least rewarding one. Choose and Avoid accuracies are calculated as the probability of choosing A when paired against C, D, E, and F, (80 vs. 70, 60, 40, 30, and 20) and the probability of choosing C, D, E, or F when they are paired with B (20 vs. 30, 40, 60, 70, and 80), respectively.

Previous patients and genetic studies have demonstrated a functional connection between these two measures and the basal ganglia pathways. For example, PD, whose indirect pathway dominates over the direct one due to SNc dopaminergic cell death, have higher Avoid accuracy than Choose accuracy. Furthermore, this pattern is reversed when drugs are administered that overcompensate the direct pathway activity. Additionally, individuals with genetic alleles that cause a greater production of dopamine receptors in the direct pathway tend to be Choosers rather than Avoiders; conversely, individuals whose alleles cause a greater number of dopamine receptors in the indirect pathway tend to be Avoiders (Frank et al., 2007).

2.2.4. Reading Skill Measure.

Because differences in performance in tasks measuring semantic processing can arise from individual differences in reading exposure, the Author Recognition Test (ART, Stanovich & West, 1989) was used as a proxy for reading skill. In the ART, participants are provided with a list of 40 real literary authors, alongside 40 foils to detect random responding. Participants are instructed to select authors they recognize, and a total score is computed based on the number of real authors selected minus foils. Although simple, the ART has been previously validated as a proxy for reading experience, and studies using eye-tracking have found that the ART is a significant predictor of time spent encoding each word (Moore & Gordon, 2015).

2.2.5. Analyses.

2.2.5.1. Behavioral Data Cleaning.

Target words in the WPT were cleaned on an individual participant basis for response time outliers, defined as trial response times greater than or lower than three standard deviations from that individual’s mean. Additionally, participants whose performance was below chance (50%, given the two-forced choice task) were eliminated from analysis.

2.2.5.2. Participant Groups.

Participant groups were created using PSS Choose and Avoid scores. Since one of the guiding assumptions of this investigation was that one's ability to learn from *both* positive and negative feedback may have an effect on semantic processing, groups were created using a relative score where Avoid was subtracted from Choose, which resulted in scores between 100 and -100. Thus, high values reflected participants who learned primarily from positive feedback (Choosers), low values reflected participants who learned primarily from negative feedback (Avoiders), and values around zero reflected individuals who learned equally as well from positive and negative feedback (Balanced). This resulted in 44 Choosers, 38 Avoiders, and 52 Balanced participants.

Table 1
PSS task descriptives by group

Group	*Mean Choose Accuracy*	*Mean Avoid Accuracy*	*Mean Choose - Avoid Accuracy*	*Mean Train Trials*
Choosers	86.79	47.02	39.77	211.40
Avoiders	44.74	85.69	-40.95	175.30
Balanced	74.02	71.94	2.08	181.20

2.2.5.3. Linear Mixed Effects Model.

The data were analyzed using linear mixed effects (LME) models, as this method has been previously shown to outperform the traditional procedures such as ANOVA (e.g., can adequately handle imbalances in group sizes; Kristensen & Hansen, 2004). LME models were specified using the R lme4 package (Bates, Mächler, Bolker, & Walker, 2014). The model was specified using the following formula:

$$\text{Target Accuracy} \sim \text{ISI} \times \text{Condition} \times \text{PSS Group} + \\ (1 + \text{Condition} \mid \text{Participant})$$

where the dependent variable is Target Word accuracy, the fixed-effects term is the factors for ISI (short or long) × Condition (dominant or subordinate) × PSS Group (Choosers, Balanced, or Avoiders), and the random effects term allows for each participant to have a different slope (or effect) for Condition, while intercepts and slopes for each participant by Condition are allowed to be correlated (e.g., higher intercepts may also have steeper slopes). An identical model was specified using Target Response Time as the dependent variable of interest. Finally, a type III ANOVA with Satterthwaite's method was used to test for significance between the factors of interest in the LME model.

2.3. Results

2.3.1. Group Level Results.

2.3.1.1. Word-Pair Task.

A repeated-measures ANOVA reveals a significant difference in Target Accuracy ($F(3, 396) = 584.30$, $p < 0.001$). Tukey's honestly significant difference (HSD) test on the orthogonal contrasts shows that differences in Target Accuracy are driven by lower accuracy to the Subordinate condition relative to all others. There were no other differences in Target Accuracy between conditions. A second repeated-measures ANOVA on Target RT also revealed statistically significant differences (F(3, 396) = 143.70, p < 0.001). Post-hoc tests revealed significant differences across all pairwise contrasts. For details, see *Table 2*.

Table 2
WPT Conditions orthogonal contrasts for Target Accuracy and RT

Measure	*Contrast*	*difference*	*SE*	*df*	*t ratio*	*p-val*
Accuracy	Dominant - Subordinate	0.38	0.011	396	33.87	< 0.001
Accuracy	Dominant - Related	-0.014	0.011	396	-1.26	0.57
Accuracy	Dominant - Unrelated	0.005	0.011	396	0.41	0.97
Accuracy	Subordinate - Related	-0.39	0.011	396	-35.13	< 0.001
Accuracy	Subordinate - Unrelated	-0.37	0.011	396	-33.46	< 0.001
Accuracy	Related - Unrelated	0.02	0.011	396	1.67	0.34
RT	Dominant - Subordinate	-171.31	11.41	396	-15.01	< 0.001
RT	Dominant - Related	35.42	11.41	396	3.10	0.01
RT	Dominant - Unrelated	-114.94	11.41	396	-10.07	< 0.001
RT	Subordinate - Related	206.73	11.41	396	18.11	< 0.001
RT	Subordinate - Unrelated	56.37	11.41	396	4.94	< 0.001
RT	Related - Unrelated	-150.36	11.41	396	-13.17	< 0.001

Notes. RT values represent milliseconds. p-value adjustment done using Tukey’s HSD method

2.3.1.2. Probabilistic Stimulus Selection Task.

Subjects performed similarly across Choose (M = 69.78, SD = 22.24) and Avoid (M = 67.99, SD = 22.22) trials. Furthermore, as in previous studies using the PSS Task (Stocco, Murray, Yamasaki, Renno, Nguyen, & Prat, 2017; Frank et al., 2007; Frank & Hutchinson, 2009), Choose and Avoid trials were not correlated (r(138) = -0.12, p = 0.14).

2.3.2. Linear Mixed Effects Analysis.

The LME model predicting Target Accuracy had a total explanatory power (conditional R^2) of 90.62%, in which the fixed effects explained 68.43% of the variance (marginal R^2). The model revealed a significant main effect of Condition (F(1, 131) = 1096.33, p < 0.0001). A significant two way interaction between Condition ✕ ISI was also observed (F(1, 262) = 6.47, p = 0.012), alongside a significant three-way interaction between Condition ✕ ISI ✕ PSS Group (F(2, 262) = 3.86, p = 0.022). Marginal two-way interactions were observed for Condition ✕ PSS Group (F(2, 131) = 2.45, p = 0.87) and also ISI ✕ PSS Group (F(2, 262) = 3.00, p = 0.51). For details, see *Figure 3*.

A follow-up analysis using the orthogonal contrasts extracted from the LME model suggest that the three-way interaction between Condition ✕ ISI ✕ PSS Group is explained by higher accuracy to Target Words during the Subordinate condition observed in PSS Choosers (difference = 0.083, t(166.85) = 2.41, p = 0.028) and Avoiders (difference = 0.086, t(166.95) = 2.60, p = 0.017), relative to the Balanced group, during the long ISI.

The LME model predicting Target RT only found a main effect of Condition (F(1, 131) = 206.42, p < 0.001). There were no significant main effects of ISI (F(1, 262) = 0.008, p = 0.93) or PSS Group (F(1, 262) = 0.67, p = 0.52), nor interactions between Condition ✕ ISI (F(1, 262) = 0.226, p = 0.63), Condition ✕ PSS Group (F(2, 262) = 0.221, p = 0.81), ISI ✕ PSS Group (F(2, 262) = 0.540, p = 0.58), or Condition ✕ ISI ✕ PSS Group (F(2, 262) = 0.61, p = 0.54).

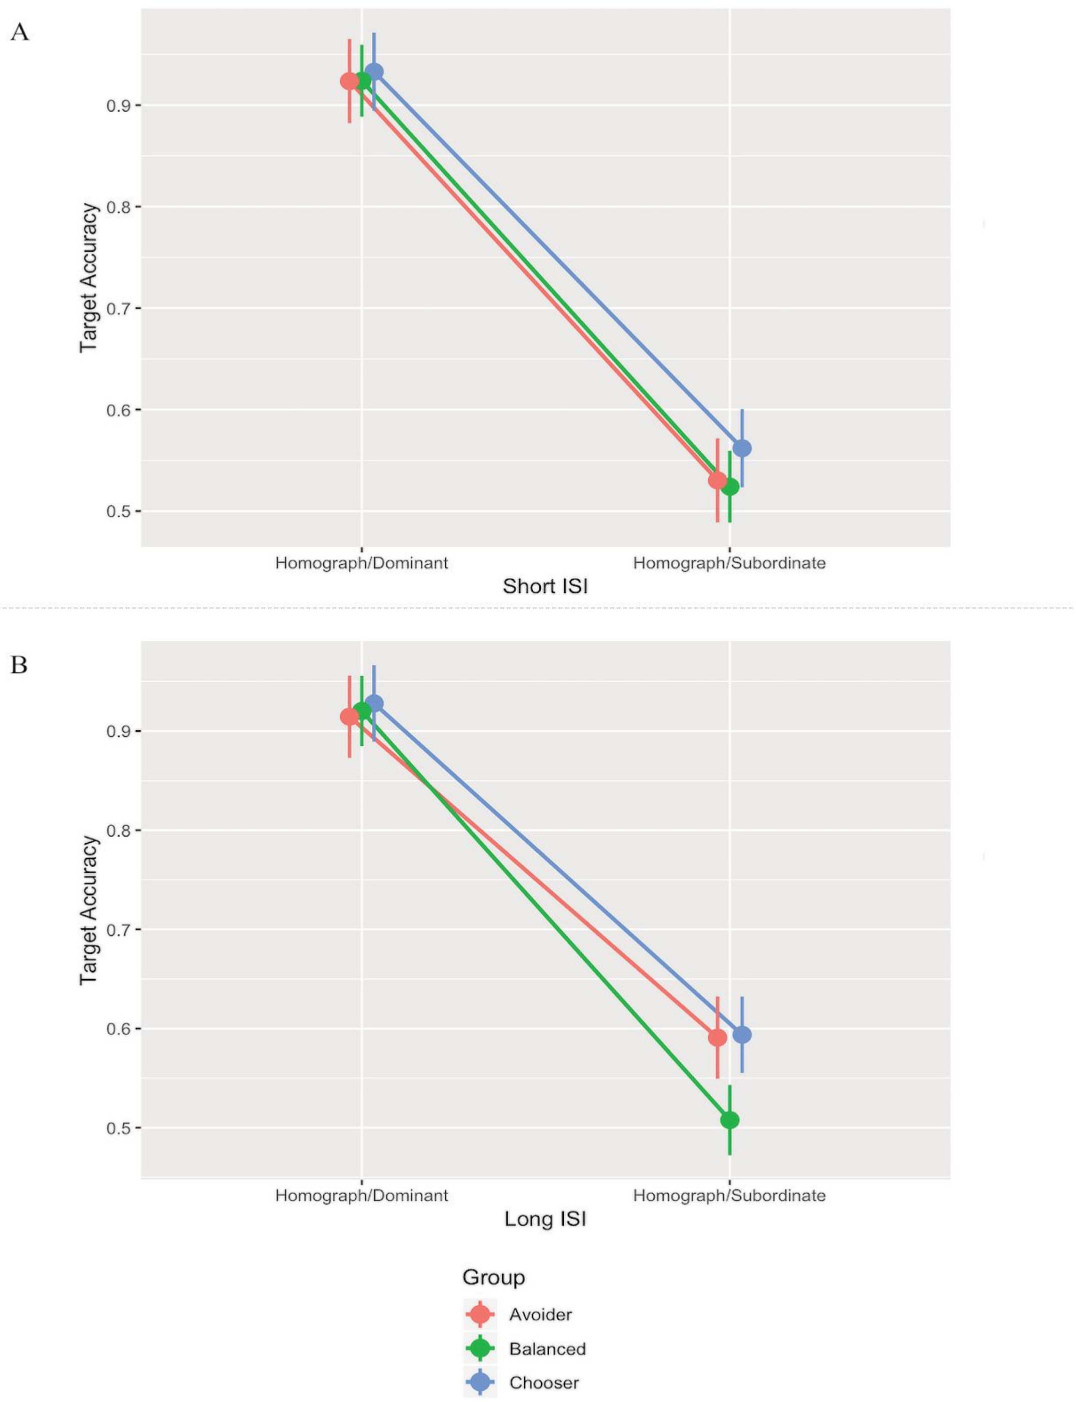

Figure 3. (A) Accuracy for conditions at the short ISI. (B) Accuracy for conditions at the long long ISI. Error bars represent a 95% CI.

2.3.3. Control Measure of Reading Experience.

The ART score differences were computed in order to ensure that differences in sensitivity to the subordinate word meaning observed was not driven by reading experience. There were no differences between Choosers and Balanced participants (t(94) = -0.44, p = 0.66), nor between Avoiders and Balanced participants (t(88) = -0.06, p = 0.95) that could account for the effect observed in the LME model results reported previously. A second LME model was estimated using ART as a covariate:

$$\text{Target Accuracy} \sim (\text{ISI} \times \text{Condition} \times \text{PSS Group}) / \text{ART} +$$
$$(1 + \text{Condition} \mid \text{Participant})$$

While the interaction of interest between Condition × ISI × PSS Group remained statistically significant under the presence of the ART covariate (F(12, 262) = 1.89, p = 0.36), this model provided a poorer model fit due to a greater number of estimated parameters (see *Table 3* for details).

Table 3
Fit indices for LME model predicting WPT Target Accuracy with and without ART covariate

Model	*Df*	*AIC*	*BIC*	*logLik*	*Dev*	*ChiSq*	*Chi Df*	*p > ChiSq*
Without ART	16	-934.3	-865.9	-.09	.02			
With ART	26	-821.6	-710.4	.08	.19	0	10	1

Notes. For AIC, a lower value is better; BIC, lower is better; Log-Likelihood, higher is better; Deviance, lower is better.

2.4. Discussion of Experiment One: Evidence for Prediction.

The current project explored the hypothesis that human linguistic ability, and specifically semantic processing, is dependent on core basal ganglia mechanisms. The results of this experiment suggest that individual differences in learning from positive or negative feedback explain variability in semantic selection processes. Interestingly, these results reproduce and extend, by artificially segmenting a continuum of basal ganglia-mediated Choose and Avoid learning in a healthy population, findings observed in clinical groups. As mentioned previously herein, Parkinson's and Huntington's Disease patients show abnormal semantic priming effects, with disrupted automatic semantic ambiguity resolution and sustained multiple competing representations. Taken together, these findings highlight the importance of a competitive dual-path reinforcement learning system that gives rise to learning from both positive and negative feedback.

Possible alternative explanations exist for the current set of experimental results. For example, many theoretical and computational models of basal ganglia functioning focus on its role as "gates" that modulate prefrontal cortex functioning through selection (or Choose) and inhibition (or Avoid) mechanisms. Thus, under this framework, we would expect to find that Choosers would manage conflict in multiple competing representations by selecting the relevant or dominant word meaning, while Avoiders would inhibit the subordinate meaning. This is, however, not what is observed in the behavioral results, where both Choosers and Avoiders show identical performance in the subordinate condition after the long delay. This pattern of results is most compatible with a reinforcement learning explanation of statistical learning, where a one-path mechanism (akin to traditional TD-learning) would overestimate the utility of the lower frequency meaning. In other words, it is possible that Choosers are overly sensitive to low-frequency reward probabilities, while Avoiders are less sensitive to high-frequency reward probabilities. This results in a misrepresentation of the relative frequency effect observed between the dominant and subordinate word meanings.

This proposed role of the basal ganglia in reinforcement learning through statistical mapping of the rich and dynamic linguistic environment, and engaging in live predictive processing may properly account for its involvement across multiple language processing modalities. As mentioned previously, work by Kotz and colleagues (Kotz et al., 2009) has presented the idea that the basal ganglia’s role in morphosyntactic processing is simply a result of these structures engaging in statistical learning and predictive processing. Specifically, this work has focused on the role of the basal ganglia in morphosyntactic processing because of the high degree of sequence-like processing involved in the parsing of grammatical information. However, this framework is compatible with semantic processing and the statistical learning and prediction processes necessary for building representations of meaning. It can also be extended to multiple processing domains, including those beyond linguistic processing (e.g., non-linguistic cognitive functioning and motor processing).

Chapter 3. A Computational Cognitive Model of Context-free Semantic Selection

Cognitive architectures allow for the creation of models of cognition by offering plausible mechanisms, obtained from decades of perceptual, cognitive, and neuroscientific research, to simulate human behavior. As discussed in a recent review of cognitive modeling methods (Gregory, Ashby & Helie, 2011), these methods provide various advantages in advancing our understanding of human cognition. For example, cognitive architectures provide additional constraints than behavioral models. Furthermore, the use of tools with well-specified cognitive (and sometimes neuroscientific) underpinnings may give rise to theories that are testable with lab experimental procedures. For these reasons, I employed a cognitive architecture to explore the mechanisms underlying individual differences in basal ganglia dual-path reinforcement learning and semantic selection processes during ambiguity resolution.

3.1. The Use of Cognitive Architectures in Cognitive Neuroscientific Research

The use of cognitive architectures as an investigative tool in cognitive research dates back to the 1970s (Newell, 1973). During its early decades, cognitive architectures remained niche in the field of cognitive psychology, but it has risen to a high degree of popularity in recent years. The reason for this rise in popularity can be understood through the advantages provided by the various types of models commonly used to study cognitive processes. The first type falls under the category of *descriptive* models, where the purpose is to provide a description of an observed outcome. Most (but not all) statistical and mathematical models are examples of descriptive models. As applied to the cognitive sciences, these kinds of models describe the observed outcome of a cognitive process quantified by psychometric tests. In other words, descriptive models are useful tools for describing the *what* aspect of neurocognitive phenomena. The second type of model belongs to the family of *process* models, and their purpose is to provide a mechanistic account of how the observed outcome was obtained. In an applied context in the cognitive sciences, these models usually describe the cognitive mechanisms (and their interactions) that gave rise to behavior previously quantified using psychometric tests. Additionally, one of the most valuable aspects of process models is that they create concrete predictions that can be tested with experiments or through additional/iterative modeling. Thus, process models provide a unique advantage and move cognitive neuroscience explorations beyond descriptions focused on the *what* or *where,* to those also focused on the *how*.

3.2. ACT-R

While dozens of cognitive architectures exist, the Adaptive Control of Thought-Rational architecture (ACT-R, henceforth; Anderson, Bothell, Byrne, Douglass, Lebiere, & Qin, 2004) is arguably the most popular to date (Kotseruba & Tsotsos, 2018). ACT-R is a general theory of cognition that enables the development of complete models capable of end-to-end simulations of a complete task. The popularity of ACT-R is driven by human models of cognition that strike a careful balance between simplicity in implementation and a high degree of cognitive and sometimes neuroscientific reality. As a result of this popularity, over 700 peer-reviewed publications contain computational cognitive models using the ACT-R cognitive architecture.

ACT-R is a computational architecture that is built from cognitive, perceptual, and neuroscience theory. The architecture makes use of declarative and procedural information storage systems to interact with the environment (e.g., behavioral task) and read from, modify, and update its internal states (see *Figure 4*. for a schematic of the ACT-R cognitive architecture). Declarative knowledge takes the form of "chunks," allowing an ACT-R model to represent facts and other static knowledge necessary for behaving in its environment. For example, a chunk in ACT-R can contain information such as “Orcas are dolphins” or even something more complex such as “Bank is a biased homograph with dominant meaning associated with money and subordinate meaning associated with river.” Procedural knowledge takes the form of a production rule in ACT-R, usually referred to simply as “productions.” Productions represent a model’s knowledge about state-action pairs, and this gives rise to the procedural steps necessary for performing a given task. For example, a production in ACT-R can contain information such as “If word on-screen is related to previous word, use ‘related’ key for manual response.” Thus, through a combination of declarative memories encoded in chunks, and state-action pairs encoded in productions, ACT-R can be used to model complex human behavior with varying levels of fidelity.

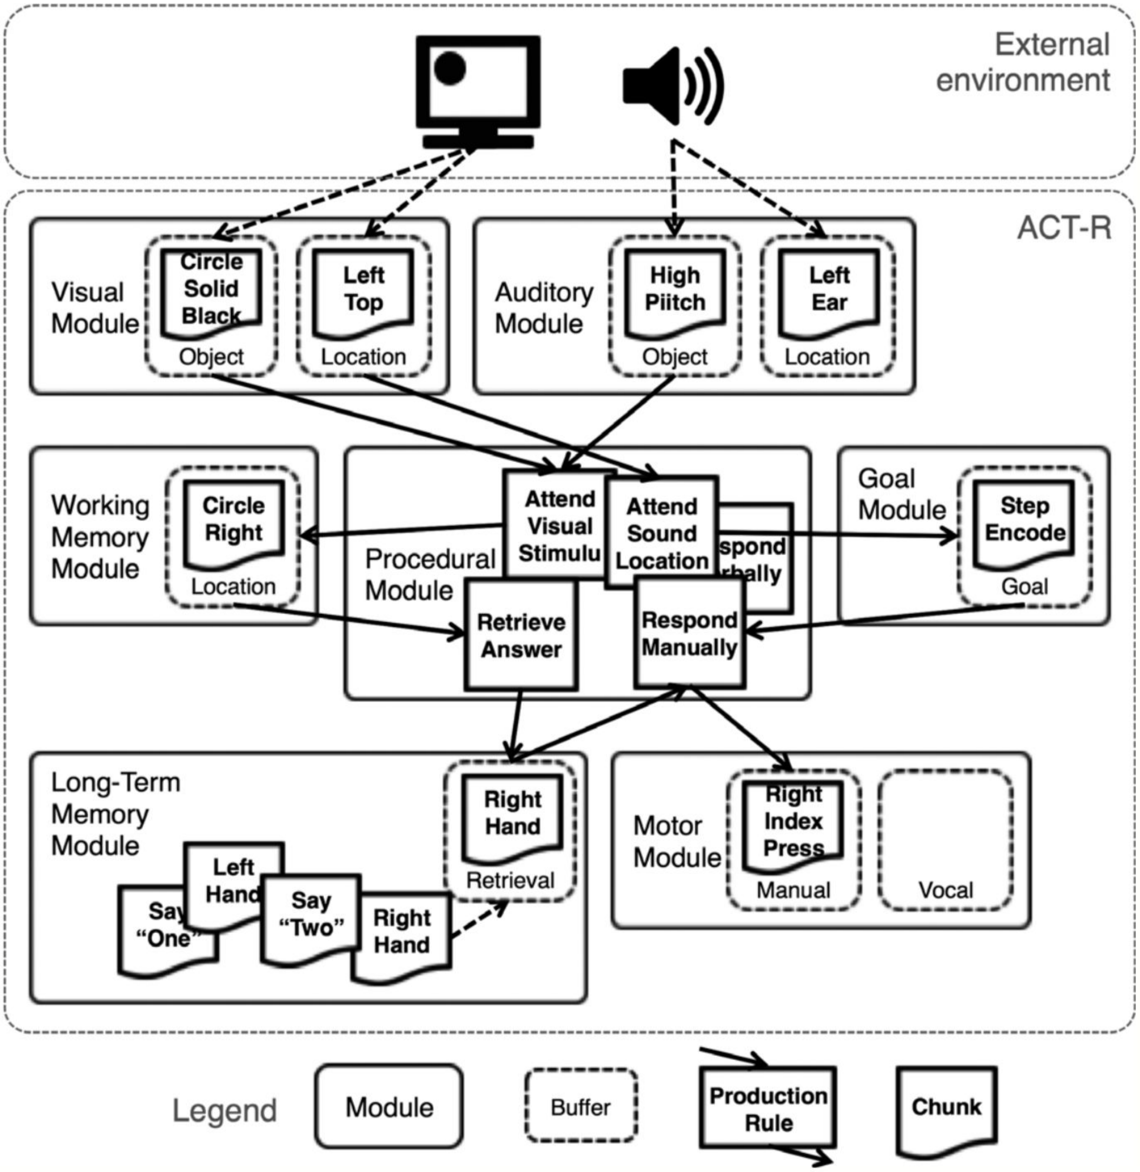

Figure 4. A schematic of ACT-R showing the relationship between the multiple components in the architecture and their interaction (reproduced from Stocco, 2018).

3.3 A Reinforcement Learning Model of Basal Ganglia Functioning

Findings involving individual differences in task performance (typically non-linguistic tasks, e.g., Simon Task) have been traditionally framed under a basal ganglia selection and inhibition framework. However, based on the observed results in Experiment One, the Word-Pair Task, I further explore the hypothesis that the signature role of basal ganglia in RL may more accurately explain its role in semantic selection. Critical for the current investigation, the activity of the basal ganglia is often modeled as reflecting Temporal Difference (TD) algorithms. However, TD-based methods do not accurately reflect the computations of the basal ganglia, which are the result of the contributions of two opposing functional pathways, namely the direct and indirect pathways. Their contribution has been modeled as the sum of competing reinforcement learning systems (Frank, Seeberger, & O'Reilly, 2004; Stocco, 2018). Individuals vary in the learning rates of the two pathways as a function of biological parameters (such as the density of dopamine receptors: Frank et al., 2007), and external factors (administration of dopamine: Frank et al., 2004). Crucially, individual differences in the preponderance of each pathway can be indirectly measured through the Probabilistic Stimulus Selection Task (PSS) task (Frank et al., 2004; Stocco et al., 2017). Thus, the current investigation tests the hypothesis that individual differences in PSS task behavioral indices of basal ganglia pathways will be related to performance in a lexical prime style task. Furthermore, I expect that a balance in functioning across both pathways is critical for optimal semantic selection and ambiguity resolution.

3.4. Methods: A Reinforcement Learning Model of Semantic Processing.

A theoretical model was implemented to examine predictions on the relationship between reward learning and semantic selection (code for the model is available on the Cognition and Cortical Dynamics Laboratory's GitHub repository: *http://github.com/uwccdl/bagels_actr*). The model was developed in the ACT-R cognitive architecture (Anderson et al., 2004).

Consistent with my hypothesis, the process of semantic selection was modeled as a procedural process; specifically, as the competition between alternative production rules. Traditionally, in ACT-R, the competition between productions is resolved through a softmax algorithm that preferentially selects the actions with the highest estimated utility, a scalar quantity that depends on the history of previous successful uses of the production and is learned through a reinforcement-learning algorithm. This type of production system in ACT-R is said to reflect basal ganglia neurocomputations (Stocco, Lebiere, O'Reilly, & Anderson, 2012), and in fact, the production module can even output an estimated BOLD signal originating from the striatum during task performance (Stocco & Anderson, 2008). However, this implementation of updating a production's utility is done making use of only positive reward, reflecting more traditional TD methods. Currently, ACT-R updates the utility U_p^t of a production *p* at time *t* according to the equation:

$$U_p^t = U_p^{t-1} + \alpha(R^t - U_p^{t-1})$$

where α is the learning rate parameter and *t* is the time at which the reward R^t is received. In the case of a traditional ACT-R model, a semantic selection process could be implemented as two productions, one that selects the dominant meaning of a word following the presentation of the ambiguous homograph, and another that selects the subordinate meaning. Depending on the outcome of this selection (e.g., selected and represented the subordinate meaning, but the interlocutor was referring to the dominant one), the production would update its utility value only when selecting for the correct meaning. As described previously in section *1.5.3. Dual-Path*

Reinforcement Learning, this is incompatible with reinforcement learning in the basal ganglia but can be solved through a modification of the ACT-R production system. Thus, the model employed herein is based on a model of the role of basal ganglia *dual-path* reinforcement learning in implicit learning tasks (Stocco, 2018). According to this model, the conflict between the two pathways can be simulated in ACT-R as a conflict between the selection of opposite and symmetric *productions*, that is, state-action pairs that implement minimal cognitive steps. Productions representing the direct (d1) pathway implements "Go" actions, while those representing the indirect (d2) pathway represent opposite "No Go" actions.

For example, the choice between two options in the PSS task, A and B, can be represented as the competition between two alternative pairs of productions, “Choose A” and “Avoid A” and “Choose B” and “Avoid A”. The modified equations for updating the utility of the Choose and Avoid productions, based on their use of positive and negative feedback are:

$$U_p^t = U_p^{t-1} + \alpha(d_1 R^t - U_p^{t-1}) \text{ for “Choose” productions}$$

$$U_p^t = U_p^{t-1} + \alpha(d_2 R^t - U_p^{t-1}) \text{ for “Avoid” productions}$$

Here, the inclusion of the d_1 and d_2 parameters allows for an individual’s natural dopamine receptor composition in the direct and indirect pathway to modulate the size of the reward that comes from learning from positive and negative feedback. Stocco (2018) provides critical evidence for why this dual-path implementation is necessary, namely that differentially altering the learning rates of “Choose” and “Avoid” productions captures (a) individual differences due to differential expressions of dopamine genes which affect the dominance of each pathway

(Frank et al., 2007), and (b) the effects of basal ganglia dopamine agonists on pathologies such as Parkinson's Disease (Frank et al., 2004). The different Choose and Avoid learning rates will be indicated as α_C and α_A, respectively.

An ambiguity resolution experiment can also be understood as a two-alternative forced choice (2AFC) task in the context of lexical retrieval. In essence, two homographs are competing for access during semantic retrieval and selection processes. Consequently, for each choice, two competing selections are performed. Thus, if the two homographs are a dominant and a subordinate interpretation of the same written word, each of them will have two production rules associated with them, "Choose Dominant" and "Avoid Dominant", and "Choose Subordinate" and "Avoid Subordinate". Contrary to traditional 2AFCs, in lexical access, the two options are not equivalent in terms of response times. Selection of the dominant meaning is usually associated with much shorter retrieval times than selection of the non-dominant meaning. In the current model, this was captured by causing those production rules that select the subordinate meaning ("Avoid Dominant" and "Choose Subordinate") to have a longer execution time. As a consequence, under short ISI, the subordinate meaning is never successfully selected. Under longer ISIs, however, participants *do* have a chance to select these meanings, so that the eventual firing of productions that select the subordinate interpretation could result in the successful retrieval of the least common meaning of the homograph. For a schematic of the ACT-R model, see *Figure 5*.

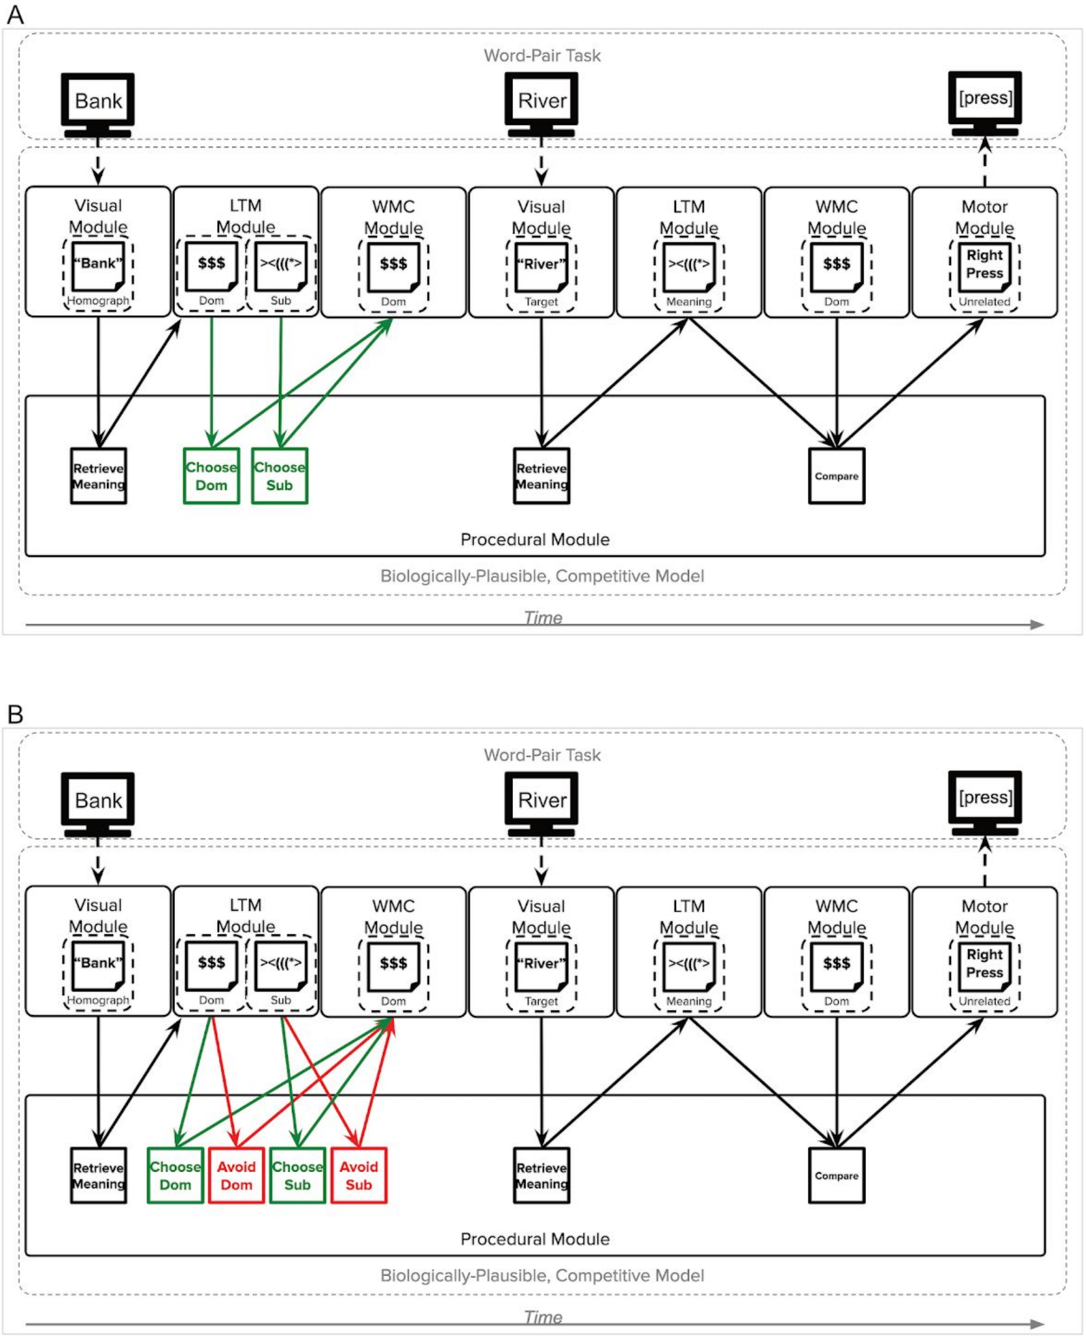

Figure 5. A schematic of the ACT-R WPT model. (A) Shows how a traditional ACT-R model, with TD-like (or single path) learning, would select between competing meaning representations.

(B) Shows how the modified model, which uses basal ganglia inspired dual-path learning, can choose *or* avoid each meaning independently.

Finally, to derive predictions from the model, I conducted an extensive set of simulations of the utility values associated to production rules under different reward conditions, corresponding to different situations in which the selection of the dominant or subordinate meaning is correct. Specifically, I examined a hypothetical situation in which the dominant meaning is contextually correct 80% of the time and the subordinate 20% of the time, much like the relative proportion of meaning frequency in the homographs used in Experiment One. To simulate the large amount of experience with the occurrence statistics of different lexical items that is associated with adult native speakers, the model was let to learn the corresponding utility values until they reached asymptotic values. In other words, the model simulates the end point of a history of linguistic experience with different probabilities of reward. Thus, performance in this model reflects adult learners that have been exposed thousands of times to ambiguous homographs and have had to perform semantic selection and later receive feedback on the basis of that selection. Furthermore, and important to the current investigation, these simulations of language experience were conducted under different learning rate parameters. The parameter values were chosen to reflect the values that were found to best capture the genetic variance of dopamine receptors in healthy adults. These are previously estimated from fitting data to the PSS task using maximum likelihood to identify the parameter space that results in the most probable values of d_1 and d_2 given task performance (Stocco, 2018). From this I was able to simulate three groups of individuals, exhibiting a preference to learn from positive feedback ($\alpha_C = 1.5 \mid \alpha_A = 1.0$), a preference to learn from negative feedback ($\alpha_C = 1.0 \mid \alpha_A = 1.5$), or no preference between the two ($\alpha_C = 1.5 \mid \alpha_A = 1.5$). These parameters are associated with different performance profiles in the PSS task, corresponding to a preference for “Choose A”, for “Avoid B”, or for a balance between the two (Stocco, 2018).

3.5. Model Predictions

To generate predictions, the model was run for 1,000 times under the different values of α_C and α_A associated with Choosers, Avoiders, or Balanced individuals. The model predicts that, for the short ISI, all three groups should perform at chance for the subordinate meaning, with no significant differences in performance. For the long ISI, however, the model predicts that Avoiders and Choosers should have a greater-than-chance performance for the subordinate condition, while Balanced individuals should still perform essentially at chance. Specifically, Avoiders will perform with an accuracy of 62%, Choosers will have an accuracy of 63%, and Balanced individuals will have an accuracy of 55%. Note that these predictions are parameter-free, and come remarkably close to the actual results of the behavioral experiment. In the model, this asymmetry in behavior is due to the fact that different initial learning rates α_C and α_A result in biased estimates of success when selecting dominant and subordinate meanings, respectively. In particular, the model predicts that Choosers would tend to overestimate the probability of the subordinate meaning, while Avoiders would tend to underestimate the probability of the dominant meaning, with both cases resulting in a tendency to favor the selection of the subordinate meaning. Under balanced learning rates, instead, the model correctly estimates the rarity of the subordinate meaning and tends to select it significantly less often.

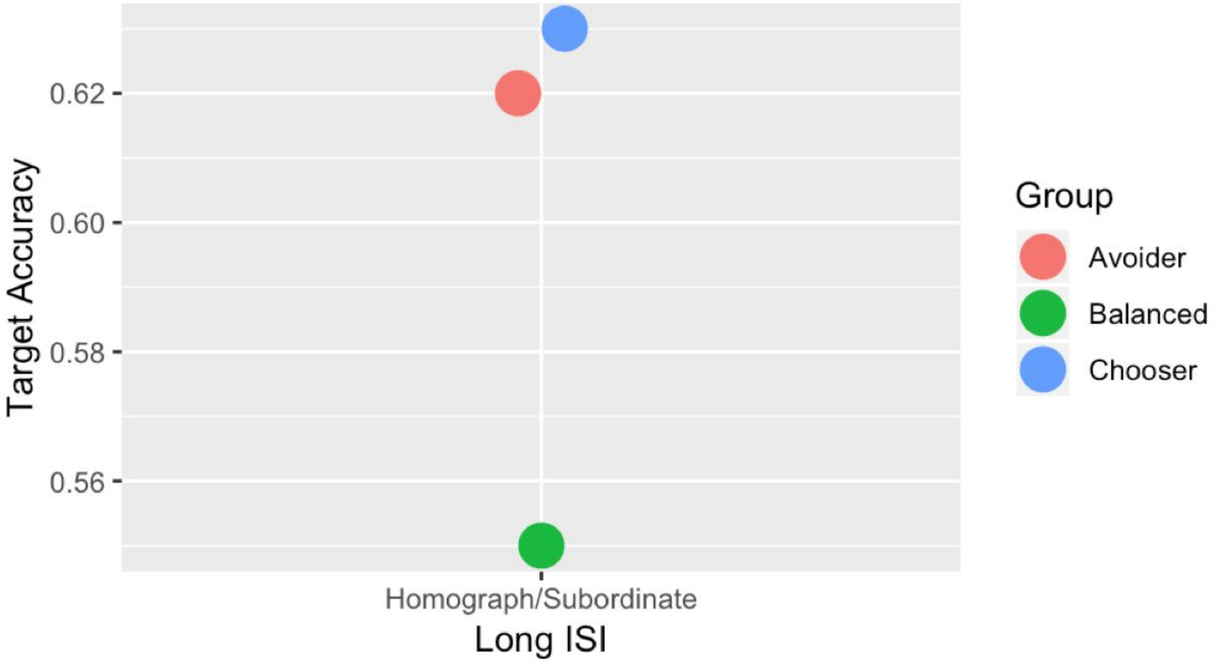

Figure 6. Model simulation predictions for the homograph subordinate long ISI condition across the three participant groups.

3.6. Computational Cognitive Model Discussion

The simulated results obtained from the computational cognitive model suggest that individual differences in basal ganglia competitive pathway dynamics give rise to variability in semantic selection processes. Particularly interesting is the prediction that, despite engaging in statistical learning through vastly different mechanisms, Choosers and Avoiders both end up over-representing the predictive utility of stimuli that occur infrequently in the environment. These results can be extended to make predictions about how individual differences in direct and indirect pathway functioning can affect other types of ambiguity processing. Specifically, operating under the assumption that dominant and subordinate meanings are represented as *relative* frequencies associated with ambiguous homographs (i.e., both frequencies sum up to one), then a higher predictive value for subordinate meanings will result in a lower value for dominant meanings. This essentially translates to Choosers and Avoiders representing the various word meanings related to biased homographs in a way that is closer to balanced homographs which occur with similar relative frequencies in the environment. Thus, a semantic ambiguity resolution task making use of balanced homographs should see similar performances for Choosers and Subordinate participants across biased and balanced homographs. Furthermore, in the case of the Word-Pair Task, this low-frequency stimulus is the subordinate meaning associated with a biased ambiguous homograph. However, the prediction of over-representing an infrequently occurring stimulus is not specific to lexical items or to the semantic ambiguity priming paradigm utilized herein. Instead, this prediction generalizes to any environmental stimulus, linguistic or not, and places this reinforcement learning and prediction account of the basal ganglia as a promising hypothesis to account for their role across various levels of linguistic processing.

Chapter 4. Experiment Two: Semantic Ambiguity Resolution In Context

4.1. Operationalization of Semantic Processing In Context

Up to this point, our understanding of the role of basal ganglia reinforcement learning in semantics stems from selection in a context-free setting. While this has provided a simple setting to build a set of predictions and test them, it is also not how people process language in everyday settings. Thus, from an ecological validity perspective, it is important to study semantic selection processes embedded in context. Furthermore, context is understood to play an important role in the processing of semantic ambiguities, as it creates constraints on the various active meanings when encountering an ambiguous word (Vitello & Rodd, 2015). Additionally, context adds complexity to the temporal dynamics of semantic ambiguity processing and modulates the extent to which frequency information drives selection. For example, ambiguous sentences that resolve ambiguity late may induce strategies that *encourage* the maintenance of multiple meanings in mind until that information has been processed. This type of interaction between meaning frequency and context is important in furthering our understanding of how basal ganglia may bias and support semantic selection. Therefore, a second experiment was used in order to measure ambiguity resolution in context and further explore the relationship between individual differences in basal ganglia functioning and semantic processing.

This experiment made use of sentences with a late disambiguating context, where an ambiguous homograph appeared early on, the sentence context remained ambiguous towards either homograph meaning and finally disambiguated towards the end of the sentence (for more details, see section *4.2.2. Semantically Ambiguous Sentences Task*). Furthermore, the experiment made use of balanced homographs in addition to biased homographs. The addition of balanced homographs enabled the testing of predictions made by the computational cognitive model, namely that Choosers and Avoiders have biased estimates of the predictive utility of selecting the subordinate meaning of a word. In other words, the relative predictive value of biased homographs will be represented as something closer to balanced homographs.

Based on the results reported in chapters two and three, the non-disambiguating words between the ambiguous homograph and the sentence-final disambiguating information will act as a prolonged delay between "critical word" and "semantic target." In experiment one, Balanced participants were found to have decreased access to the subordinate meaning following a long word-pair delay of 850 ms. The computational cognitive model provided causal evidence for this effect arising from Choosers and Avoiders representing dominant and subordinate word meanings as being closer to relatively equal frequencies. Therefore, compared to Balanced participants, Choosers and Avoiders are expected to show smaller ambiguity effects (as indexed by increases in reading speed) upon encountering the word that disambiguates towards the subordinate meaning of a biased homograph. Furthermore, for Choosers and Avoiders, response times to the disambiguating region in the biased subordinate condition should resemble response times to the disambiguating word region in the balanced condition.

4.2. Methods

4.2.1. Participants.

Informed consent was obtained from participants prior to the experiment, as required by the Institutional Review Board at the University of Washington. All participants were recruited using the Psychology Department Participant Research Pool. Additionally, participants were compensated with course credit for their participation in this experiment. Data were collected from 200 healthy participants (125 females, mean age = 19.0 years). Sixteen participants were removed from analyses due to issues with data collection or poor task performance, defined as performance below chance. All participants completed the experimental tasks in four Latin-squared orders to control for possible order or fatigue effects induced by the sentence task and PSS task length.

Additionally, the Language Experience and Proficiency Questionnaire (LEAP-Q; Marian, Blumenfeld, & Kaushanskaya, 2007) was used to obtain indices of English language experience, as unlike in the WPT, this experiment did not make use of only monolingual participants. The reason for this is that during the year of data collection for this experiment, the participants available through the Psychology Research Pool were largely bilingual. This, alongside study exclusion criteria (e.g., could not have previously participated in the WPT experiment as they share ambiguous homographs), led to a very low sample size (n = 55) after multiple quarters of data collection. Thus, participants were allowed to sign up irrespective of language profile. Instead, I used LEAP-Q measures of English Speaking and Understanding Proficiency for to create a composite score of English Proficiency. Composite scores were then used as a covariate in statistical analyses to account for the effects of variability in English fluency on semantic ambiguity processing.

4.2.2. Semantically Ambiguous Sentences Task.

A self-paced reading task (Ferreira & Henderson, 1990), was designed to measure the effects of semantic ambiguity within the context of a sentence (Semantically Ambiguous Sentences or SAS task, hereafter). Each trial consisted of a self-paced reading sentence, followed by a relatedness judgment used to ensure participants were carefully reading the sentences and engaging with the task. For the relatedness judgment, participants were asked to button press if word presented was related or unrelated to the previously shown sentence, and had up to 3 seconds to respond. For the self-paced reading sentences, the experiment made use of a serial demasking style paradigm where a fully masked sentence appeared at the onset of the trial in the center of the screen (with the same number of characters to the right and left of the screen center). The mask was generated with a continuous underline, and each word was masked individually. A monospaced, or fixed-width font was used in order to ensure that the length of each underline accurately reflected the width of each word. Participants pressed the “m” key in the keyboard with their right index finger in order to begin the serial demasking procedure. Each key press caused the demasking “window” to move from left-to-right, leading to only one word being de-masked at any given point in time during the reading of the sentence. This “moving window” was implemented in order to prevent eye regressions during reading, which would have confounding effects on the self-paced reading values for each word position. For an example of the press-by-press demasking procedure used, see Figure 7 below.

Sentences were constructed using stimuli obtained from the norming study by Twilley and colleagues (Twilley et al., 1994). Specifically, ambiguous homographs were selected as the critical word and appeared early in the sentence, in word positions two or three. In the example provided in *Figure 7*, the ambiguous homograph is in word position number two (see “boxer” in *Figure 7*, keypress 2). The sentences were constructed so that there were at a minimum three additional words following the critical item before any disambiguating information appeared, This allowed room for spillover of reaction time effects following the presentation of the critical word. The distance between critical and disambiguating word led to the disambiguating word occurring in positions seven, eight, or nine. In the example provided in Figure 7, the disambiguating word occurs in word position seven (see “vet” in *Figure 7*, keypress 7). Similarly, in order to allow room for spillover of reaction time effects following the disambiguating information, there was a minimum of three words following the disambiguating word.

Keypress	Sentence Stimulus On Screen
	___ _____ ___ _______ _____ ___ ___ ____ ___ ______
1	The _____ ___ _______ _____ ___ ___ ____ ___ ______
2	___ boxer ___ _______ _____ ___ ___ ____ ___ ______
3	___ _____ was _______ _____ ___ ___ ____ ___ ______
4	___ _____ ___ hurting _____ ___ ___ ____ ___ ______
5	___ _____ ___ _______ after ___ ___ ____ ___ ______
6	___ _____ ___ _______ _____ the ___ ____ ___ ______
7	___ _____ ___ _______ _____ ___ vet ____ ___ ______
8	___ _____ ___ _______ _____ ___ ___ gave ___ ______
9	___ _____ ___ _______ _____ ___ ___ ____ him ______
10	___ _____ ___ _______ _____ ___ ___ ____ ___ shots.

Figure 7. A sample of the self-paced reading serial demasking stimuli. Each row shows what the participant sees on-screen after the “m” key is pressed.

The SSA task made use of two types of ambiguous homographs. The first type was balanced homographs, defined as a word with two possible meanings that occur with equal relative frequency in the world (for details, see Twilley et al., 1997). An example of a balanced homograph is the word “bat,” which can be related to the sport or the flying mammal. The second type of homograph was biased homographs (or imbalanced homographs), defined as a word with two possible meanings where one occurred with a much higher frequency than the other. This definition of biased homographs is the same as used in the context-free semantic processing study (for details, see section *2.2.2. Context-free Semantic Ambiguity Task,* or Twilley et al., 1997). An example of a biased homograph is the word "bank," where the dominant or high-frequency meaning is related to the financial establishments, while the subordinate or low-frequency meaning is related to the slopes bordering a river.

Sentences were constructed so that until the disambiguating word, either word meaning was equally possible. Furthermore, sentences that disambiguated to each of two meanings of an ambiguous homograph were created, alongside a control sentence that was identical to each of the sentences containing an ambiguous homograph. Thus, for each homograph, a total of four sentences were constructed, two experimental sentences for each meaning and two control sentences. Participants were presented with one version of the two sentences corresponding to each ambiguous homograph, alongside the control sentence for the *other* ambiguous homograph meaning. Therefore, two stimulus lists were created and participants were presented only one of two lists. The following examples illustrate these manipulations across one set of balanced homographs and another set of biased homographs:

Balanced Homograph:

Critical Sentence - Meaning One (List A):

The / seal / was / hidden / so / the / plumber / had / to / work / carefully.

Control Sentence (List B):

The / damage / was / hidden / so / the / plumber / had / to / work / carefully.

Critical Sentence - Meaning Two (List B):

The / seal / was / visible / as / it / swam / through / the / murky / water.

Control Sentence (List A):

The / fish / was / visible / as / it / swam / through / the / murky / water.

Biased Homograph:

Critical Sentence - Dominant Meaning (List A):

The / cabinet / was / ready / for / the / meeting / with / the / new / president.

Control Sentence (List B):

The / secretary / was / ready / for / the / meeting / with / the / new / president.

Critical Sentence - Subordinate Meaning (List B):

The / cabinet / was / finished / and / the / carpenter / could / finally / sell / it.

Control Sentence (List A):

The / shelf / was / finished / and / the / carpenter / could / finally / sell / it.

Each list contained a total of 88 sentences: 20 sentences with biased homographs, 12 sentences with biased homographs that disambiguated towards the dominant meaning, 12 sentences with biased homographs that disambiguated towards the subordinate meaning, and 44 unambiguous control sentences. In order to make sure participants understood the task, four practice sentences were provided during the instruction phase of the task. Additionally, prior to the presentation of the 88 experimental and control sentences, two dummy "start" sentences were provided in order to allow participants to develop a steady self-paced reading speed. Finally, reaction times were recorded for each individual word presented during the self-paced reading portion of the trial, while reaction times and response accuracies were recorded for the relatedness judgment probe following each sentence.

Prior to data collection, all sentences were normed online using Amazon's Mechanical Turk. Fifty L1 speakers of American English read all experimental and control sentences and rated each for comprehension on a scale of one to four, where 1 = cannot be understood and 4 = can be easily understood. Prior to norming, an average of 2 was selected as the minimum acceptable score, and any sentences not meeting this threshold would be rewritten. The final set of self-paced reading items received an average score of 3.49, with a standard deviation of 0.31.

4.2.3. Probabilistic Stimulus Selection Task.

The Probabilistic Stimulus Selection task used in this experiment is identical to the one used in experiment No. 1, the context-free semantic ambiguity investigation. Please refer to section *2.2.3 Probabilistic Stimulus Selection Task* for details on this task.

4.2.4. Reading Skill Measure.

The control measure for reading skill, the Author Recognition Test used in this experiment is identical to the one used in experiment No. 1, the context-free semantic ambiguity investigation. Please refer to section *2.2.4 Reading Skill Measure* for details on this task.

4.2.5. Analyses.

4.2.5.1. Behavioral Data Cleaning.

SSA task data was cleaned for response time outliers. Specifically, self-paced reading response times to each word were cleaned on a by-subject basis. Individual word response times that were 3 times larger or smaller than the participant mean were discarded from analysis. Furthermore, participants with performance below chance, or 50%, for the relatedness judgment probe were discarded from analysis. Additionally, two dependent variables of interest were extracted from the self-paced reading data. The first was participants' response time to the critical word, aggregated by participant over word positions zero through three (where word position zero is the ambiguous homograph). The second was participants' response time to the disambiguating word, aggregated by participant over word positions zero through three (where word position zero is the first word providing disambiguating information).

4.2.5.2. Participant Groups.

Similar to the procedure used for creating participant groups in Experiment One, groups were created using PSS Choose and Avoid scores. A relative score was computed, where Avoid was subtracted from Choose, which resulted in scores between 100 and -100. Thus, high values reflected participants who learned primarily from positive feedback (Choosers), low values reflected participants who learned primarily from negative feedback (Avoiders), and values centered around zero reflected individuals who learned equally as well from positive and negative feedback (Balanced). This grouping resulted in 53 Choosers, 59 Avoiders, and 72 Balanced participants.

Table 4

PSS task descriptives by group

Group	*Mean Choose Accuracy*	*Mean Avoid Accuracy*	*Mean Choose - Avoid Accuracy*	*Mean Train Trials*
Choosers	87.97	53.77	34.20	167.50
Avoiders	54.87	83.47	-28.60	176.90
Balanced	72.22	70.49	1.73	195.00

4.2.5.3. Linear Mixed Effects Model.

The SSA task data were analyzed using linear mixed effects (LME) models. Like in the analysis used in Experiment One, this method was chosen as it has been previously shown to outperform the traditional procedures such as ANOVA (Kristensen & Hansen, 2004) and can adequately handle imbalances in group sizes (e.g., PSS Groups). For validation purposes, the same results were reproduced using ANOVA (although not reported herein). LME models were specified using the R lme4 package (Bates et al., 2014). Two separate models were specified for each of the dependent variables of interest, response times to the critical word and response times to the disambiguating word.

The first model tested for reaction time differences at the onset of the critical word across conditions, with response times averaged over word positions 1-3 following the critical word onset. This model was specified using the following formula:

Critical Window RT ~
(Condition ✕ PSS Group) / English Proficiency +
(1 | Participant)

where the dependent variable is response time to the critical word, the fixed-effects term is the factors for Condition (biased-dominant, biased-subordinate, or balanced) ✕ PSS Group (Choosers, Balanced, or Avoiders), with English Proficiency as the composite score covariate. The random effects term allows for each participant to have a different intercept for word response time. A random effects structure containing variation by condition for each participant was attempted but the model failed to converge due to negative model eigenvalues. A type III ANOVA with Satterthwaite's method was used to test for significance between the factors of interest in the critical word LME model.

The second model tested for reaction time differences at the onset of the disambiguating word across conditions, with response times averaged over word positions 1-3 following the disambiguating word onset. This model was specified using the following formula:

Disambiguating Window RT ~

(Condition ✕ PSS Group) / English Proficiency +

(1 | Participant)

where the dependent variable is response time to the disambiguating word, the fixed-effects term is the factors for Condition (biased-dominant, biased-subordinate, or balanced) ✕ PSS Group (Choosers, Balanced, or Avoiders), with English Proficiency as the composite covariate. The random effects term allows for each participant to have a different intercept for word response time. A random effects structure containing variation by condition for each participant was attempted but the model failed to converge due to negative model eigenvalues. A type III ANOVA with Satterthwaite's method was used to test for significance between the factors of interest in the critical word LME model.

4.3. Results

4.3.1. Group Level Results.

4.3.1.1. Semantically Ambiguous Sentences Task.

A mean response time (controlled by reading speed) window for the critical word was calculated by averaging the critical word at its onset and the following three words. A mean response time window was also calculated for the disambiguating word window using this same method. A one by three ANOVA for the three experimental conditions showed no significant effects of condition over reading speed at the critical word window (F(2, 549) = 0.034, p = 0.97). However, and as expected, a one by three ANOVA for the three experimental conditions showed statistically significant effects of condition over reading speed at the disambiguating word window (F(2, 549) = 20.78, p < 0.001). Relative to the Biased-Dominant condition, Balanced sentences were not significantly different (t(183) = -0.05, p = 0.96), while Biased-Subordinate sentences were (t(183) = 5.56, p < 0.001). Relative to Biased-Dominant trials, the reading speed of Biased-Subordinate trials was 17.16 ms slower (beyond baseline, as RTs here are controlled for reading speed) in the averaged window, suggesting a high ambiguity effect upon encountering the sentence disambiguating context for this condition type.

4.3.1.2. Probabilistic Stimulus Selection Task.

Subjects performed similarly across Choose (M = 70.56, SD = 20.75) and Avoid (M = 69.06, SD = 19.71) trials. Furthermore, as in Experiment One (see section *2.3.1.2. Probabilistic Stimulus Selection Task* for details) and in previous studies using the PSS Task (Stocco et al., 2017; Frank et al., 2007; Frank & Hutchinson, 2009), Choose and Avoid trials were not correlated (r(184) = 0.038, p = 0.082).

4.3.2. Linear Mixed Effects Analysis.

The LME model predicting reading time during the critical word window revealed a marginal main effect of PSS Group (F(2, 178) = 2.98, p = 0.053). A marginal interaction between Condition ✕ PSS Group was also observed under the effect of the English Proficiency covariate (F(9, 266.75) = 1.74, p = 0.079). For further details see *Figure 8A*. A second LME model predicting reading time during the disambiguating word window revealed a significant main effect of PSS Group (F(2, 178) = 4.97, p = 0.008). Additionally, a significant interaction between Condition ✕ PSS Group was found (F(4, 356) = 2.55, p = 0.039), which remained significant under the effects of the English Proficiency covariate (F(9, 266.75) = 2.46, p = 0.010). A follow-up analysis using the orthogonal contrasts extracted from the LME model for disambiguating word position suggest that the interaction between Condition × PSS Group is explained by slower response times to the disambiguating word for the Biased-Subordinate condition observed the PSS Balanced group (t(530.18) = -3.288, p = 0.001), suggesting a high ambiguity effect upon encountering the sentence disambiguating context for this condition type. For further details, see *Figure 8B*.

4.3.3. Control Measure of Reading Experience.

The ART score differences were computed in order to ensure that differences in sensitivity across conditions was not driven by reading experience. There were no differences between Choosers and Balanced participants (t(184) = 0.76, p = 0.45), nor between Avoiders and Balanced participants (t184) = 1.00, p = 0.32) that could account for the effects observed in the LME models results reported previously.

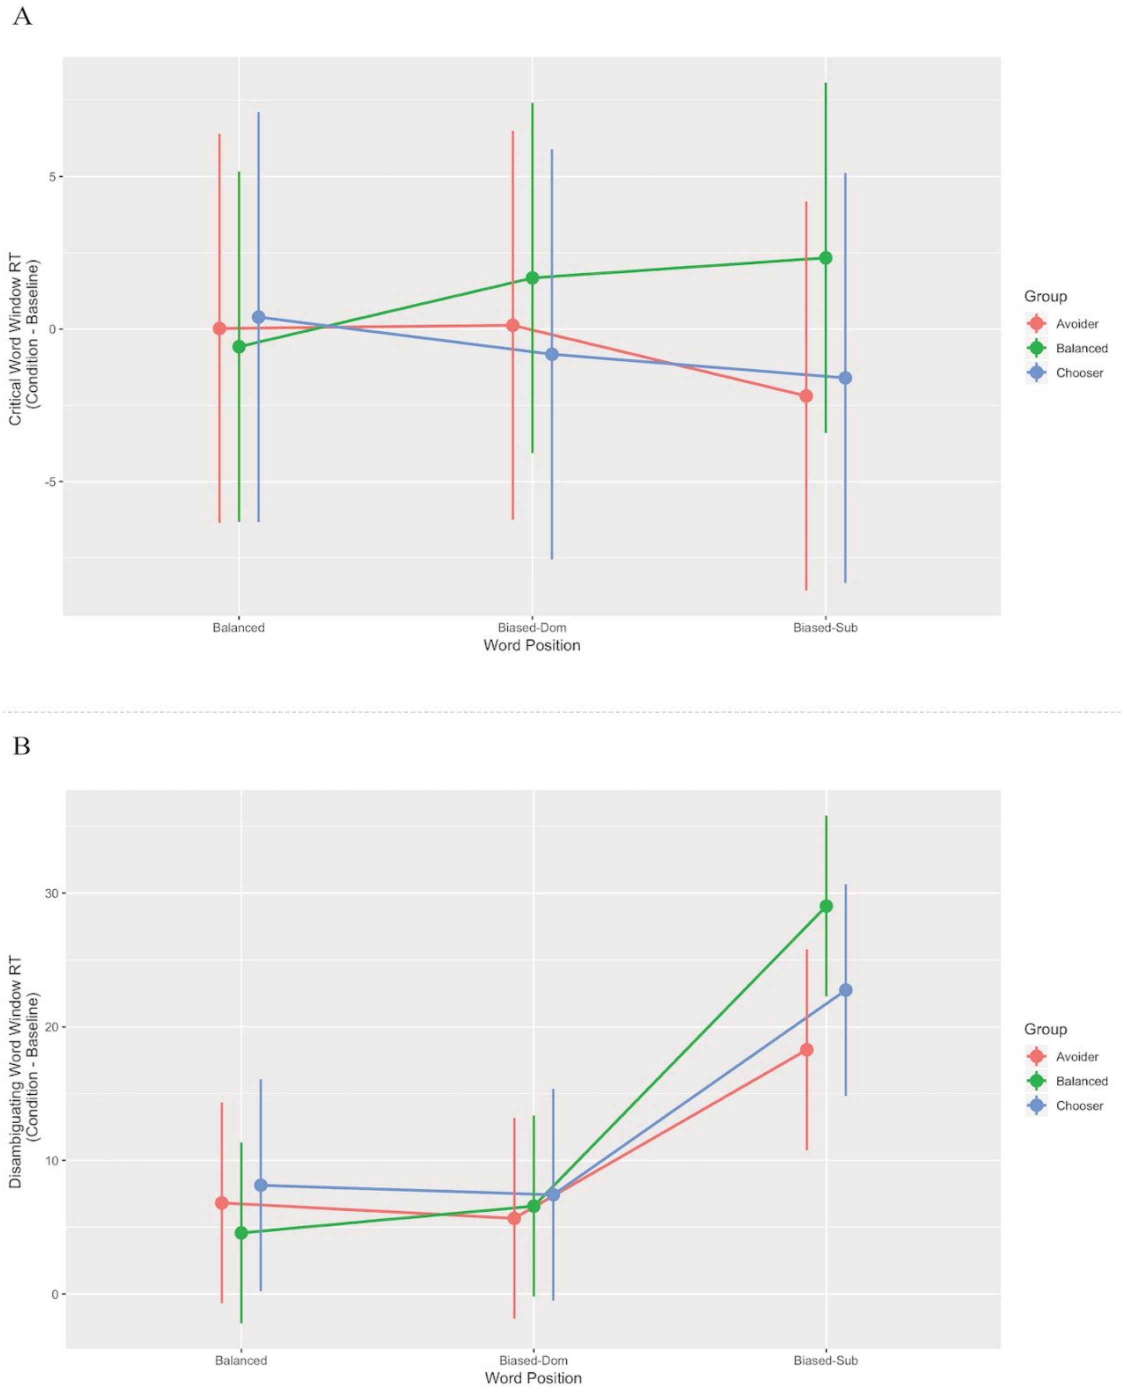

Figure 8. Response times to the (A) critical word window and (B) disambiguating word window across conditions. Error bars represent a 95% confidence interval.

4.4. Discussion of Experiment Two

The current experiment aimed to test the hypothesis that the role of the basal ganglia in semantic ambiguity resolution in context can be explained through reinforcement learning action selection mechanisms. The results of this experiment extend results previously reported in *Chapters 2 and 3*, namely that the basal ganglia and their role in learning and prediction via reinforcement provide a robust neurocomputational framework to account for individual differences in semantic selection.

This self-paced reading paradigm shows that reading speed for ambiguous sentences that disambiguate to a low-frequency meaning varies as a result of individual differences in basal ganglia pathway functioning. Specifically, while the three participant groups showed slow-downs in reading the disambiguating region of the Biased-Subordinate sentences, this effect was only statistically significant for the Balanced participant group (see *Figure 8B* for details). The slow-down in reading speed provides details about how (im)probable the Biased Subordinate sentence ending was for readers. In the case of Choosers and Avoiders, their minimal slow-down suggests that the subordinate sentence ending was probable and remained available from the presentation of the ambiguous homograph through the disambiguating region. This made it less surprising and resulted in Choosers and Avoiders having non-significant slow-downs in reading speed. For Balanced participants, on the other hand, disambiguation to the subordinate meaning induces a larger slow-down, suggesting that this subordinate sentence ending was improbable and no longer available. This is in line with an account that Balanced participants do not select for the Biased-Subordinate meaning amongst competing ones given an accurate mapping of the low probability of occurrence.

Finally, it is possible that the results reported previously are driven by a late effect observed in the Choosers and Avoiders that may be masked in averaging over the selected window. In other words, the averaged time window may reduce sensitivity in detecting differences in effects over time. An additional LME model fitted using disambiguating word position as a factor (0-3, see below).

Disambiguating Region RT ~

(Condition ✕ PSS Group X Word Position) / English Proficiency +

(1 | Participant)

In this expanded model, a significant interaction between Condition ✕ PSS Group X Word Position was observed under the effects of the English proficiency covariate (F(36, 1063.4) = 1.7347, p = 0.005). The results show that Choosers and Avoiders do not show significant ambiguity effects to the Biased-Subordinate condition, relative to the Biased-Dominant one (see *Figure 9*), even with three words for spillover following the onset of the disambiguating information.

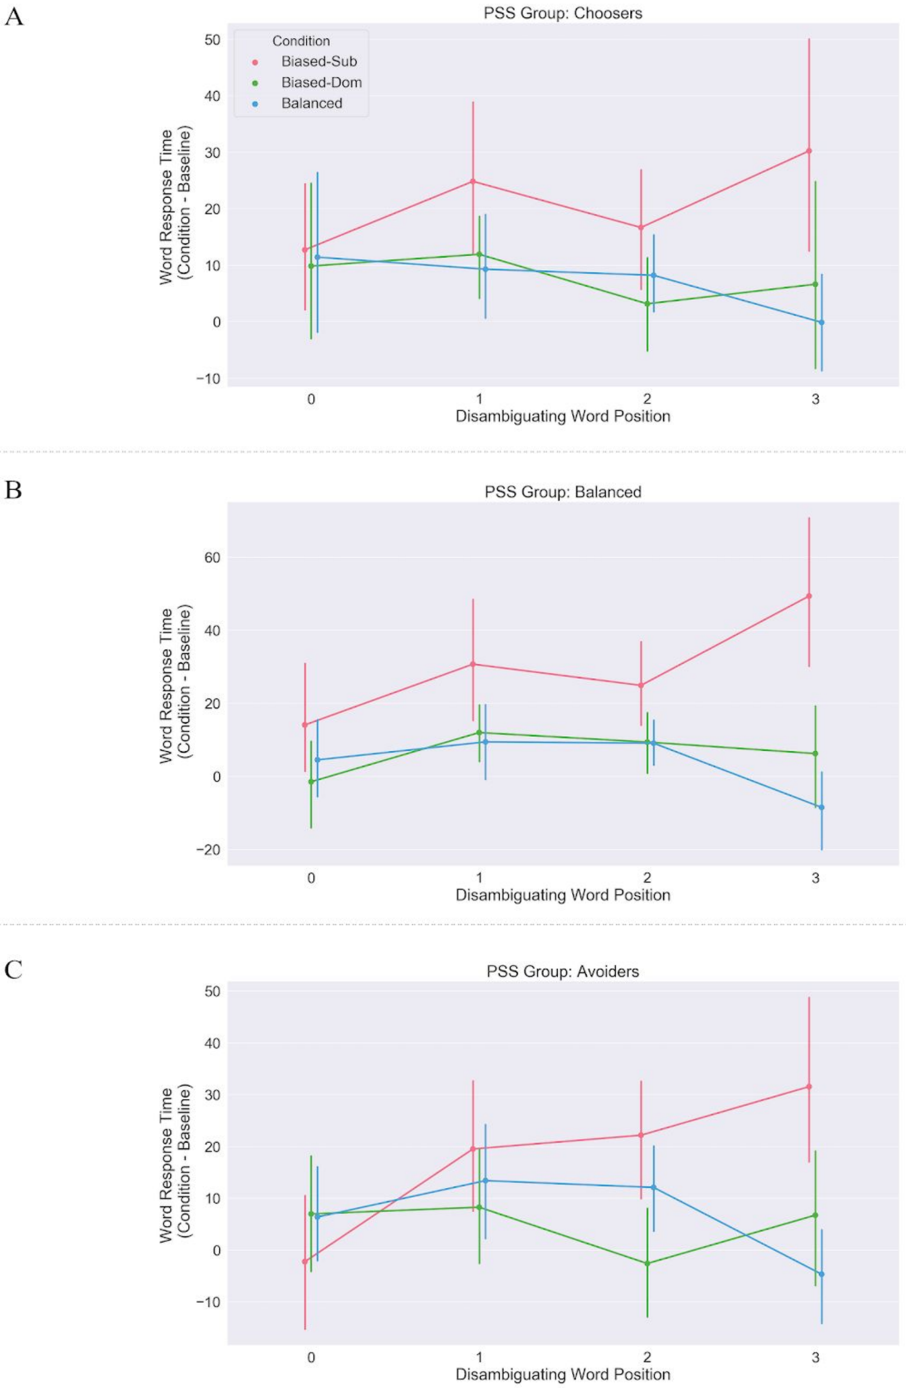

Figure 9. Response times from the disambiguating word (at position zero) by condition for (A) Choosers, (B) Balanced, and (C) Avoiders. Error bars represent a 95% CI.

Chapter 5. Discussion

5.1. Basal Ganglia Reinforcement Learning Influences on Semantic Selection

The results of the computational model simulations and two experiments used herein converge to support the novel hypothesis that **competitive reinforcement learning dynamics in the basal ganglia drive individual differences in semantic selection processes**. These novel findings have critical implications both for models of the neurobiology of semantic selection and for understanding the nature of individual differences in word knowledge. Each of these implications is discussed subsequently, along with their broader theoretical importance.

5.1.1. Dual Pathway Reinforcement Dynamics in Semantic Selection.

This dissertation constitutes the first series of investigations relating dual-path learning dynamics to semantic ambiguity resolution. The findings reported herein quantify the impacts of learning from positive and negative feedback, based on the biological system under investigation (the basal ganglia nuclei). Importantly, across two experiments and one computational model, I have demonstrated that behavioral estimates of the relative strength of the direct and indirect pathways of the basal ganglia relate to individual differences in semantic ambiguity resolution processes. Importantly, this research shows that learning from mistakes, or omissions, is equally important as learning from reward for the way in which an individual represents the likelihood of reward of different linguistic outcomes.

Additional evidence for this perspective stems from a natural language processing (NLP) model of that was very recently published on distributional semantics for learning word meanings from contextual co-occurrence patterns (Johns, Mewhort, & Jones, 2019). This model used both positive *and* negative information to learn word meanings from surrounding contextual information. Importantly, like the balanced participants in the current group who make equal use of both positive and negative information during learning and prediction, the distributional semantics neural network model outperformed similar models making use of only positive (or negative) information to select for specific word meanings. The positive and negative information used by this model is obtained from the error signal derived during prediction, much like a subcortical RPE. And although the researchers do not mention the basal ganglia (their focus is not to build biologically plausible models after all), the mechanisms implemented come remarkably close to the reinforcement learning neurocomputations that the basal ganglia give rise to via its competitive pathway dynamics. In summary, the results reported herein provide a biologically based account for how dual-pathway reinforcement learning shapes prediction processes in language.

5.1.2. Implications for Individual Differences in Reading.

The results described herein also have critical implications for our understanding of individual differences in semantic selection, or lexical representation more broadly. Crucially, this research demonstrates that an individual's estimated utility of selecting a specific meaning is driven not only by their degree of experience with that meaning and its alternatives, but also by an interaction between their biology, which weights positive and negative evidence individually, and their experience. The importance of fast and accurate access to lexical representations for reading more broadly is perhaps best illustrated by Perfetti's lexical quality hypothesis (Perfetti, 2007). However, models of skilled reading such as this emphasize the importance of experience in building accurate representations of how likely some event will occur based on a "database" of relevant previous experiences.

To illustrate this, when a Balanced participant reads the word “bank” they co-activate the associated “money” and “river” related meanings. Selection happens in line with their reinforcement-learning-based estimate that the word “bank” much more frequently relates to the “money” sense of the word than to the “river” sense of the word and this subordinate meaning is unavailable for the semantic relatedness judgment, resulting in poor task performance (for this condition, only). On the other hand, when an Avoider reads the word “bank,” they are much less likely to select for the high-frequency meaning as they have a strong expectation that “river” may occur based on biased estimates of success. Thus, the biasing signal provided by the basal ganglia during semantic selection can be seen as reflecting an individual’s estimated representation of the relative frequency of a meaning associated with a lexical form. Critically, individual differences in learning from positive and negative reward drive this estimation process. This is a significant departure from the way ambiguity resolution processes are discussed in the literature, generally with little regards to factors contributing to individual variability in performance. As typically discussed, individuals encountering a sentence that disambiguates towards a very low-frequency meaning will experience a large ambiguity effect, and any variability observed in this process is attributed to experiential factors. For example, a lexical priming study showed that a recent encounter with a homograph in context will bias the subsequent interpretation of that word, even after a 20+ minute long delay (Rodd, Cutrin, Kirsch, Millar, & Davis, 2013). The results of this study were interpreted to reflect experience- or recency-driven adaptations of lexical representations. My research, instead, suggests that these effects are likely driven by encountering an ambiguous word in context, which provides an individual with an additional instance of learning from predicting a meaning outcome from a specific lexical form.

5.1.3. Semantic Selection in Context.

Psycholinguistic and neurobiological theories of semantic selection and ambiguity resolution have posited that meaning frequency is the strongest contributing factor in ambiguity resolution (Vitello & Rodd, 2015). Thus, while context is known to play a number of roles in ambiguity processing, the relative frequency with which a meaning occurs in the context of an ambiguous word is the primary driver of the behavioral and neural responses observed. The Semantically Ambiguous Sentences (SAS) task provided further insight about *what* the basal ganglia biasing signal reflects during semantic selection, as context adds complexity to the temporal dynamics of semantic selection. Specifically, a sentence that starts out ambiguously and contains disambiguating information near the end may induce a cognitive strategy of maintaining multiple meanings in mind until enough information is obtained to engage in semantic selection (Just & Carpenter, 1992).

Based on previously outlined computational models and theories of basal ganglia in dynamically routing information and working memory capacity (WMC), it is plausible that under this type of context, basal ganglia signaling to PFC supports the maintenance of multiple meanings. This explanation is not most closely in line with the data, which showed that Balanced individuals continued to show large ambiguity effects even in the delayed disambiguation condition. They did, however, maintain both Balanced homograph meanings and did not engage in early selection, as they do not experience ambiguity effects for this condition relative to a dominant disambiguating context. Thus, Balanced learners were engaging in selection based on how likely they expected to encounter a particular word meaning. For them, maintenance or selection of a Biased-Subordinate meaning may ultimately prove disadvantageous given how infrequently these meanings occur in the real world.

5.1.4. Alternate Interpretations of Findings.

One might argue that the conclusions of this dissertation work can be reinterpreted as reflecting constraints in working memory capacity (WMC) by Balanced individuals. In the case of the Word-Pair Task, it is possible that Choosers and Avoiders have greater WMC and thus can better maintain multiple meanings, leading to higher performance accuracy in the Subordinate condition (e.g., Just & Carpenter, 1992). The same pattern of results is observed for the Sentence Task, where Choosers and Avoiders have smaller ambiguity effects to the Subordinate condition. However, the sentence task made use of an additional condition that included Balanced homographs, with relatively equal probability of occurrence. The fact that Balanced participants perform equally to Choosers and Avoiders when encountering a balanced homograph suggests that differences in the ability to maintain multiple representations does not account for the differences between groups. Instead, I argue that the availability of a word meaning associated with an ambiguous homograph is related to the estimated likelihood of its occurrence. And while I cannot definitely rule that the results observed are the outcome of differences in WMC, this account can be viewed from a slightly different perspective; it is possible that, because Choosers and Avoiders more frequently co-activate multiple representations due to biased utility estimates, they train their working memory systems over time to develop greater capacities.

Alternatively, what we measure when administering participants with WMC tasks may actually be differences in strategies for dealing with multiple competing or interfering cortical representations. This interpretation, while seemingly speculative, is in fact compatible with a line of research on how accounting for and measuring strategies affects WMC (Dunlosky & Kane, 2007). For example, when provided with explicit strategies to be used in WMC tasks, participants showed a significant improvement in performance relative to controls who received the same amount of practice on the task but without the strategy training (McNamara & Scott, 2001). Thus, it is plausible that WMC is not the driver, but instead is the result, of individual differences in selecting between multiple competing representations.

Finally, additional research provides support for why Choosers and Avoiders behave similarly despite having opposing functional pathway mechanisms. The role of dopamine in the basal ganglia has traditionally been related to that of learning. However, many computational models and patient studies have also shown that this neurotransmitter also serves the purpose of contrast enhancer. In other words, the presence of dopamine facilitates the processing of nuanced contextual information, as it modulates striatal firing of direct and indirect pathway functioning (Houk, 1997), with the end result of having this firing be more selective. Importantly, this research has shown that both, too little (e.g., PD; Whiting, Copland, & Anguin, 2005) or too much dopamine (e.g., high dose of DA agonists; Swerdlow et al., 2005), results in an overall decrease in sensitivity to contextual information. Thus, it is possible to create the analogous case for Choosers behaving as having too much dopamine and Avoiders too little. Indeed, the results observed herein can be interpreted in light of reduced sensitivity to context, where Choosers and Avoiders do not contextually discriminate between the nuances in occurrence probabilities for Biased-Dominant, Biased-Subordinate, and Balanced homographs, and therefore do not show significant effects across these three conditions. Balanced participants, on the other hand, may benefit from being at the "sweet spot" in this bell curve, with enough direct and indirect pathway functioning to show sensitivity to ambiguous words with varying degrees probability of occurrence.

5.2. Behavioral Implications of Experimental Findings

It is important to consider the reasons why a balance in direct and indirect pathway functioning manifests itself in the behavioral pattern observed for Balanced participants. Research on reading ability and WMC has traditionally framed smaller ambiguity effects as advantageous, so why would Balanced participants behave in ways that cause reinterpretations that may be cognitively costly? One might argue that the more advantageous behavior is to always process and maintain all possible meanings in parallel to avoid costly reinterpretations. A possible explanation for this may stem from research into how people construct interpretations during reading. Vast individual differences exist in the level of detail and coherence used to build meaning representations during reading. Readers may build rich representations early on in some contexts, while other times they build representations that are "good enough" (Ferreira, Bailey, & Ferraro, 2016). During the building of deep representations in reading, sentences containing ambiguous pronouns are processed more slowly and larger ambiguity effects are shown (Stewart, Holler, & Kidd, 2007). This is due to the fact that readers engaging in deep processing commit more quickly to a particular (usually high probability) meaning associated with an ambiguous linguistic stimulus. However, all readers eventually need to build a coherent understanding of the content read, and research suggests that "good enough" representations are only good enough for a short while and trigger a delay in interpretation. In the case of Balanced participants, it possible that through a better representation of the low probability of occurrence of the subordinate meaning they are able to engage in deeper processing upon encountering the Biased homograph. This means that they only have to engage in reinterpreting the ambiguous sentence a low proportion of the time, as opposed to consistently having to engage in delayed interpretation as Choosers and Avoiders may.

5.3. A General Basal Ganglia Mechanisms for Language Processing

This exploration began with the idea that focusing on a core basal ganglia neurocomputation may be a promising approach to explaining their involvement in language. While this work specifically focused on lexico-semantics as a model system, the reinforcement learning neurocomputations identified have the potential to explain basal ganglia involvement across various levels of linguistic processing. Psycholinguistic theory informs us that ambiguity can arise in syntax (MacDonald, 1994), morphology (Järvikivi, Pyykkönen, & Niemi, 2009), and even phonology (Lukatela, Frost, & Turvey, 1999). Importantly, these cases of ambiguity processing reflect a more general need for selection of linguistic information between multiple competing representations. Furthermore, the neurobiological mechanisms for morphosyntactic and phonological processing rely on inferior frontal structures immediately adjacent to (and oftentimes overlapping with) those for semantic selection. This should not come as a surprise to many, as in fact, the LIFG has most traditionally been discussed as a language region responsible for engaging is grammatical and phonological processing (Caplan & Waters, 1998; Broca, 1861; Nixon, Lazarova, Hodinott-Hill, Gough, & Passingham, 2004). Thus, existing unifying accounts suggest the LIFG plays a crucial role in general conflict resolution (Novick, Kan, Trueswell, & Thompson-Schill, 2009). Specifically, under this account LIFG involvement occurs when there is a dominant but inappropriate response or when there are multiple responses with similar strengths of activation. This general conflict resolution role of the LIFG is compatible with the interpretations derived from the current dissertation work, namely that the basal ganglia reinforcement learning neurocomputations provide biasing signals necessary for cortical selection processes.

5.4. Limitations

There are several limitations inherent in the data presented herein, as well as in my interpretations of them. First, the series of studies presented herein rely on behavioral indices and computational models to infer and simulate individual differences in basal ganglia contributions to semantic selection. Furthermore, this work aims to provide a model of the basal ganglia pathways in language but exclude the subthalamic nuclei's hyperdirect pathway entirely. This decision was made deliberately for two reasons: (1) There is currently no consensus on the existence of behavioral measures of the hyperdirect pathway that *only* measure hyperdirect pathway functioning and do so in a valid and reliable way, and (2) Functional imaging of the subthalamic nucleus is at best, difficult (with 7T MRI), and at worst largely unreliable (with 3T MRI; Forstmann et al., 2017). The impact of this decision is that we continue to have an incomplete understanding of the role of this third and important basal ganglia functional pathway in semantic processing and, more broadly, the neurobiology of language. Given the role of the hyperdirect pathway in "breaking" or overriding the output of the striatum's direct and indirect pathways, it is possible that incorporating this pathway in into models of language processing could account for additional important individual differences. Indeed, studies on the effects of deep-brain stimulators to the subthalamic nucleus of PD patients have shown that modulating hyperdirect pathway signaling may have positive (Vonberg, Ehlen, Fromm, Kühn, & Klostermann, 2016) and negative (Ehlen et al., 2017) effects on lexical processes.

Additionally, the choice of ACT-R as the computational modeling framework for this dissertation imposes constraints on the fidelity of the basal ganglia model. While ACT-R was chosen because of the versatility it offers in modeling language processes (especially when compared to additional frameworks, such as Nengo; Stewart, Bekolay, & Eliasmith, 2012), cognitive architectures are generally less well-suited for capturing complex biological properties that may influence cognitive processes of interest. For example, in modeling the direct and indirect pathway signaling dynamics, this work could have incorporated cytoarchitectonic details of basal ganglia functioning such as patch and matrix compartments, striatal cholinergic signaling, and even calcium interactions with dopamine during learning. Because the field's understanding of this level of subcortical fidelity in non-linguistic cognition remains rudimentary, more work is necessary before it can be successfully implemented into computational models of language.

Finally, the conclusions drawn from experiment two with regards to differences in depth of processing by Balanced participants are limited by the lack of indices of reading ability. Measures of reading comprehension, such as the Nelson Denny Reading Test (Brown, 1960), could provide further evidence for or against this interpretation.

5.5. Broader Impacts

The work showcased herein presents evidence for an important role of the basal ganglia in semantic ambiguity resolution. Results show that basal ganglia reinforcement learning processes can indeed account for vast individual differences in sensitivity to high and low-frequency word meanings associated with ambiguous homographs. Furthermore, this work provides a strong argument for the presence of a dual-path reinforcement learning neural system, as implemented by the subcortical basal ganglia. This work can also be used to better understand theories of the basal ganglia in other levels of linguistic processing. In everyday language use, whether it be the processing of phonemes during speech comprehension or the parsing of grammatical structures during reading, there are innumerable instances of ambiguity and the need for selection of a specific representation amongst competing ones. So while semantics was used as the model system to understand individual differences in basal ganglia functioning, the reinforcement learning framework employed by this work offers an exciting opportunity to continue exploring holistic models of language processing that incorporate contributions of subcortical structures.

Appendix

Supplementary Table 1. Studies reporting basal ganglia involvement in language (reproduced from Ceballos, Zeitlin, Prat, & Stocco, in prep)

APA Citation	Prod. / Comp.	Ling. Process	Population	Mono/Bi-lingual	Research Method	Caudate	Putamen	Globus Pallidus (GPe/GPi)	Subthalamic Nucleus (STN)	Substantia Nigra (SNr/SNc)
Abdullaev & Melnichuk (1997) *	comprehension	phonology/phonetics	Parkinson's	monolingual	DBS	x				
Abdullaev & Melnichuk (1997) *	comprehension	lexico-semantics	Parkinson's	monolingual	DBS	x				
Abutalebi et al. (2008)	production	lexico-semantics	healthy	bilingual	fMRI	x				
Abutalebi, Della Rosa, Gonzaga, Keim, Costa, & Perani (2013)	production	lexico-semantics	healthy	bilingual	fMRI		x (L)			
Abutalebi, Miozzo, & Cappa (2000)	production	lexico-semantics	BG lesion	bilingual	behavioral	x (L)				
Adrover-Roig, Galparsoro-Izagirre, Marcotte, Ferré, Wilson, & Inés Ansaldo (2011)	production	lexico-semantics	BG lesion	bilingual	behavioral	x (L)	x (L)	x (L)		
Aglioti, Beltramello, Girardi, & Fabbro (1996) *	production	lexico-semantics	BG lesion	bilingual	behavioral		x (L)			
Aglioti, Beltramello, Girardi, & Fabbro (1996) *	production	syntax	BG lesion	bilingual	behavioral		x (L)			
Aglioti & Fabbro (1993)	production	phonology/phonetics	BG lesion	bilingual	behavioral		x (L)			
Booth, Wood, Lu, Houk, & Bitan (2007)	comprehension	phonology/phonetics	healthy	monolingual	fMRI		x (L)			
Bradley, King, & Hernandez (2013) *	comprehension	lexico-semantics	healthy	bilingual	fMRI		x (L)			
Bradley, King, & Hernandez (2013) *	comprehension	lexico-semantics	healthy	monolingual	fMRI	x				
Brunner, Kornhuber, Seemüller, Suger, & Wallesch (1982) *	production	phonology/phonetics	BG lesion	monolingual	behavioral	x (L)	x (L)	x (L)		
Brunner, Kornhuber, Seemüller, Suger, & Wallesch (1982) *	production	lexico-semantics	BG lesion	monolingual	behavioral	x (L)	x (L)	x (L)		
Callan, Tajima, Callan, Kubo, Masaki, & Akahane-Yamada (2003)	comprehension	phonology/phonetics	healthy	bilingual	fMRI	x	x	x	x	
Castner, Chenery, Copland, Coyne, Sinclair, & Silburn (2007)	comprehension	lexico-semantics	Parkinson's	monolingual	DBS				x	
Chenery, Copland, & Murdoch (2002)	comprehension	lexico-semantics	BG lesion	monolingual	behavioral	x (L)	x (L)	x (L)		
Copland (2003)	comprehension	lexico-semantics	BG lesion	monolingual	behavioral	x	x	x		
Copland, Chenery, & Murdoch (2000)	comprehension	lexico-semantics	BG lesion	monolingual	behavioral	x (L)	x (L)	x (L)		
Crinion et al. (2006)	comprehension	lexico-semantics	healthy	bilingual	fMRI	x (L)				
Crosson et al. (2003) *	production	phonology/phonetics	healthy	monolingual	fMRI	x	x (R)			
Crosson et al. (2003) *	production	lexico-semantics	healthy	monolingual	fMRI	x	x (R)			

APA Citation	Prod. / Comp.	Ling. Process	Population	Mono/Bi-lingual	Research Method	Caudate	Putamen	Globus Pallidus (GPe/GPi)	Subthalamic Nucleus (STN)	Substantia Nigra (SNr/SNc)
de Diego Balaguer, Couette, Dolbeau, Dürr, Youssov, & Bachoud-Levi (2008) *	comprehension	lexico-semantics	Huntington's	monolingual	behavioral	x				
de Diego Balaguer, Couette, Dolbeau, Dürr, Youssov, & Bachoud-Levi (2008) *	comprehension	syntax	Huntington's	monolingual	behavioral	x				
Dromey, Kumar, Lang, & Lozano (2000)	production	phonology/phonetics	Parkinson's	monolingual	DBS				x	
Fabbro, Clarici, & Bava (1996) *	production	phonology/phonetics	BG lesion	monolingual	behavioral	x	x	x		
Fabbro, Clarici, & Bava (1996) *	production	morphosyntax	BG lesion	monolingual	behavioral	x	x	x		
Fabbro, Clarici, & Bava (1996) *	production	lexico-semantics	BG lesion	monolingual	behavioral	x	x	x		
Fridriksson, Ryalls, Rorden, Morgan, George, & Baylis (2007)	production	phonology/phonetics	BG lesion	monolingual	fMRI		x (L)			
Friederici, von Cramon, & Kotz (1999) *	comprehension	lexico-semantics	BG lesion	monolingual	ERP	x	x	x		
Friederici, von Cramon, & Kotz (1999) *	comprehension	syntax	BG lesion	monolingual	ERP	x	x	x		
Frisch, Kotz, von Cramon, & Friederici (2003)	comprehension	syntax	BG lesion	monolingual	ERP	x	x	x		
Garnier, Lamalle, & Sato, (2013)	production/comprehension	phonology/phonetics	healthy	monolingual	fMRI	x (R)	x			x (R)
Gentil, Chauvin, Pinto, Pollak, & Benabid (2001)	production	phonology/phonetics	Parkinson's	monolingual	DBS				x	
Gentil, Pinto, Pollak, & Benabid (2003)	production	phonology/phonetics	Parkinson's	monolingual	DBS				x	
Gil Robles, Gatignol, Capelle, Mitchell, & Duffau (2005)	production	phonology/phonetics	BG lesion	monolingual	DBS; resection	x (L)	x (L)			
Giraud et al. (2008)	production	phonology/phonetics	Persistent Developmental Stuttering	monolingual	fMRI	x				x
Gurd, Bessell, Bladon, & Bamford (1988)	production	phonology/phonetics	BG lesion	monolingual	behavioral	x(L)	x(L)	x(L)		
Henry, Berman, Nagarajan, Mukherjee, & Berger (2004)	production	lexico-semantics	BG lesion	monolingual	DTI		x (L)			
Hervais-Adelman, Moser-Mercer, Michel, & Golestani (2014)	production/comprehension	lexico-semantics	healthy	mutilingual	fMRI	x	x	x	x	
Ketteler, Kastrau, Vohn, & Huber (2008)	comprehension	lexico-semantics	healthy	monolingual	fMRI	x	x (R)			
Kotz, Frisch, & von Cramon (2003)	comprehension	syntax	BG lesion	monolingual	ERP	x (L)	x (L)	x (L)		
Kotz & Schmidt-Kassow (2015)	comprehension	syntax	BG lesion	monolingual	ERP	x	x	x		

APA Citation	Prod. / Comp.	Ling. Process	Population	Mono/Bi-lingual	Research Method	Caudate	Putamen	Globus Pallidus (GPe/GPi)	Subthalamic Nucleus (STN)	Substantia Nigra (SNr/SNc)
Kuchinke, Hofmann, Jacobs, Frühholz, Tamm, & Herrmann (2011)	comprehension	lexico-semantics	healthy	monolingual	fMRI	x	x			
Kurowski, Blumstein, & Alexander (1996)	production	phonology/phonetics	BG lesion	monolingual	behavioral	x (L)	x (L)			
Liégeois, Baldeweg, Connelly, Gadian, Mishkin, & Vargha-Khadem (2003)	production	lexico-semantics	KE family	monolingual	fMRI	x	x (L)			
Lu et al. (2010)	production	lexico-semantics	Persistent Developmental Stuttering	monolingual	fMRI		x			
Mason & Just (2007)	comprehension	lexico-semantics	healthy	monolingual	fMRI	x	x			
Moretti et al. (2003) *	production	phonology/phonetics	Parkinson's	monolingual	DBS				x	
Moretti et al. (2003) *	production	lexico-semantics	Parkinson's	monolingual	DBS				x	
Moro, Tettamanti, Perani, Donati, Cappa, & Fazio (2001)	comprehension	syntax	healthy	monolingual	PET	x (L)				
Müller, Rothermel, Behen, Muzik, Chakraborty, & Chugani (1999)	production	phonology/phonetics	BG lesion	monolingual	PET	x (R)				
Nota & Honda (2004)	production	phonology/phonetics	healthy	monolingual	fMRI		x			x
Pickett, Kuniholm, Protopapas, Friedman, & Lieberman (1998) *	production	phonology/phonetics	BG lesion	monolingual	behavioral	x	x			
Pickett, Kuniholm, Protopapas, Friedman, & Lieberman (1998) *	comprehension	syntax	BG lesion	monolingual	behavioral	x	x			
Prat & Just (2011)	comprehension	syntax	healthy	monolingual	fMRI	x (R)	x (L)			
Prat, Keller, & Just (2007)	comprehension	lexico-semantics	healthy	monolingual	fMRI	x (R)				
Radanovic & Scaff (2003)	production	phonology/phonetics	BG lesion	monolingual	behavioral	x	x	x		
Riecker, Mathiak, Wildgruber, Erb, & Hertrich (2005)	production	phonology/phonetics	healthy	monolingual	fMRI	x (L)	x	x (L)		
Seo, Stocco, & Prat (2018)	comprehension	morphosyntax	healthy	bilingual	fMRI	x				
Simmonds, Leech, Iverson, & Wise (2014)	production	phonology/phonetics	healthy	monolingual	fMRI	x	x			
Simmonds, Wise, Dhanjal, & Leech (2011) *	production	lexico-semantics	healthy	bilingual	fMRI	x		x		
Simmonds, Wise, Dhanjal, & Leech (2011) *	production	syntax	healthy	bilingual	fMRI	x		x		
Speedie, Wertman, Ta'ir, & Heilman (1993)	production/comprehension	lexico-semantics	BG lesion	bilingual	behavioral	x (R)				
Stowe, Paans, Wijers, & Zwarts (2004)	comprehension	syntax	healthy	monolingual	PET	x (R)				
Tan et al. (2011)	comprehension	lexico-semantics	healthy	bilingual	fMRI	x (L)				
Teichmann, Dupoux, Kouider, & Bachoud-Lévi (2006) *	comprehension	morphosyntax	Huntington's	monolingual	behavioral	x				

APA Citation	Prod. / Comp.	Ling. Process	Population	Mono/Bi-lingual	Research Method	Caudate	Putamen	Globus Pallidus (GPe/GPi)	Subthalamic Nucleus (STN)	Substantia Nigra (SNr/SNc)
Teichmann, Dupoux, Kouider, & Bachoud-Lévi (2006) *	comprehension	lexico-semantics	Huntington's	monolingual	behavioral	x				
Teichmann et al. (2015) *	production/comprehension	morphosyntax	BG lesion	monolingual	fMRI					
Teichmann et al. (2015) *	comprehension	syntax	BG lesion	monolingual	fMRI					
Tettamanti et al. (2005)	comprehension	phonology/phonetics	healthy	monolingual	PET	x (L)	x (L)			
Thames et al. (2012)	production	phonology/phonetics	HIV patients	monolingual	fMRI	x (L)	x (L)			
Tricomi, Delgado, McCandliss,, McClelland, & Fiez (2006)	comprehension	phonology/phonetics	healthy	bilingual	fMRI	x				
Vargha-Khadem et al. (1998) *	production	phonology/phonetics	KE family	monolingual	PET, MRI	x	x	x		
Vargha-Khadem et al. (1998) *	production	lexico-semantics	KE family	monolingual	PET, MRI	x	x	x		
Wallesch, Henriksen, Kornhuber, & Paulson (1985) *	production	phonology/phonetics	healthy	monolingual	SPECT	x	x (L)	x (L)		
Wallesch, Henriksen, Kornhuber, & Paulson (1985) *	production	syntax	healthy	monolingual	SPECT	x		x		
Wallesch, Henriksen, Kornhuber, & Paulson (1985) *	production	lexico-semantics	healthy	monolingual	SPECT	x		x		
Wallesch, Kornhuber, Brunner, Kunz, Hollerbach, & Suger (1983) *	production	syntax	BG lesion	monolingual	behavioral	x	x	x		
Wallesch, Kornhuber, Brunner, Kunz, Hollerbach, & Suger (1983) *	production	phonology/phonetics	BG lesion	monolingual	behavioral	x	x	x		
Wallesch, Kornhuber, Brunner, Kunz, Hollerbach, & Suger (1983) *	production	lexico-semantics	BG lesion	monolingual	behavioral	x	x	x		
Wartenburger, Heekeren, Abutalebi, Cappa, Villringer, & Perani (2003)	comprehension	morphosyntax	healthy	bilingual	fMRI	x (L)	x			
Watkins et al. (2002)	production	phonology/phonetics	KE family	monolingual	MRI	x				
Whelan, Murdoch, Theodoros, Hall, & Silburn (2003)	comprehension	lexico-semantics	Parkinson's	monolingual	DBS				x	
Wildgruber, Ackermann, & Grodd (2001)	production	phonology/phonetics	healthy	monolingual	fMRI		x (L)			
Zou, Ding, Abutalebi, Shu, & Peng (2012)	production	lexico-semantics	healthy	bilingual	fMRI	x (L)				

References

Abutalebi, J., Rosa, P. A. D., Castro Gonzaga, A. K., Keim, R., Costa, A., & Perani, D. (2013). The role of the left putamen in multilingual language production. *Brain and Language*, *125*(3), 307–315. http://doi.org/10.1016/j.bandl.2012.03.009

Alexander, G. E., & Crutcher, M. D. (1990). Functional architecture of basal ganglia circuits: neural substrates of parallel processing. *Trends in Neurosciences*, *13*(7), 266–271. http://doi.org/10.1016/0166-2236(90)90107-L

Alexander, M. P., Naeser, M. A., & PALUMBO, C. L. (1987). Correlations of subcorticalct lesion sites and aphasia profiles. Brain, 110(4), 961-988.

Anderson, J. R., Bothell, D., Byrne, M. D., Douglass, S., Lebiere, C., & Qin, Y. (2004). An Integrated Theory of the Mind. *Psychological Review*, *111*(4), 1036–1060. http://doi.org/10.1037/0033-295X.111.4.1036

Angwin, A. J., Chenery, H. J., Copland, D. A., Murdoch, B. E., & Silburn, P. A. (2007). The speed of lexical activation is altered in Parkinson's disease. *Journal of Clinical and Experimental Neuropsychology*, *29*(1), 73–85. http://doi.org/10.1080/13803390500507188

Arnott, W. L., Copland, D. A., Chenery, H. J., Murdoch, B. E., Silburn, P. A., & Angwin, A. J. (2011). The Influence of Dopamine on Automatic and Controlled Semantic Activation in Parkinson's Disease. *Parkinson's Disease*, *2011*(9), 1–9. http://doi.org/10.4061/2011/157072

Bates, D., Mächler, M., Bolker, B., & Walker, S. (2014). Fitting linear mixed-effects models using lme4. arXiv preprint arXiv:1406.5823.

Booth, J. R., Wood, L., Lu, D., Houk, J. C., & Bitan, T. (2007). The role of the basal ganglia and cerebellum in language processing. *Brain Research*, *1133*, 136–144. http://doi.org/10.1016/j.brainres.2006.11.074

Bradley, K. A. L., King, K. E., & Hernandez, A. E. (2013). Language experience differentiates prefrontal and subcortical activation of the cognitive control network in novel word learning. *NeuroImage*, *67*, 101–110. http://doi.org/10.1016/j.neuroimage.2012.11.018

Broca, P. (1861). Perte de la parole, ramollissement chronique et destruction partielle du lobe antérieur gauche du cerveau. Bull Soc Anthropol, 2(1), 235-238.

Brown, J. I. (1960). The Nelson-Denny Reading Test.

Caplan, D., Alpert, N., & Waters, G. (1998). Effects of syntactic structure and propositional number on patterns of regional cerebral blood flow. Journal of Cognitive Neuroscience, 10(4), 541-552.

Castner, J. E., Chenery, H. J., Copland, D. A., Coyne, T. J., Sinclair, F., & Silburn, P. A. (2007). Semantic and affective priming as a function of stimulation of the subthalamic nucleus in Parkinson's disease. *Brain*, *130*(5), 1395–1407. http://doi.org/10.1093/brain/awm059

Chenery, H. J., Angwin, A. J., & Copland, D. A. (2008). The basal ganglia circuits, dopamine, and ambiguous word processing: A neurobiological account of priming studies in Parkinson's disease. *Journal of the International Neuropsychological Society*, *14*(03), 351–364. http://doi.org/10.1017/S1355617708080491

Chenery, H. J., Copland, D. A., & Murdoch, B. E. (2002). Complex language functions and subcortical mechanisms: Evidence from Huntington's disease and patients with non-thalamic subcortical lesions. *International Journal of Language & Communication Disorders*, *37*(4), 459–474. http://doi.org/10.1080/1368282021000007730

Cohen, M. X., & Frank, M. J. (2009). Neurocomputational models of basal ganglia function in learning, memory and choice. *Behavioural Brain Research*, *199*(1), 141–156. http://doi.org/10.1016/j.bbr.2008.09.029

Copland, D. A. (2003). The basal ganglia and semantic engagement: Potential insights from semantic priming in individuals with subcortical vascular lesions, Parkinson's disease, and cortical lesions. *Journal of the International Neuropsychological Society*, *9*(07), 1041–1052. http://doi.org/10.1017/S1355617703970081

Copland, D. A., Bell, P., McMahon, K., Dissanayaka, N., Mellick, G., & Silburn, P. (2012). Subcortical Activity during Verbal Selection and Suppression in Parkinson's Disease. *Procedia-Social and …*, *61*, 50–51. http://doi.org/10.1016/j.sbspro.2012.10.074

Copland, D. A., Chenery, H. J., & Murdoch, B. E. (2000). Understanding Ambiguous Words in Biased Sentences: Evidence of Transient Contextual Effects in Individuals with Nonthalamic Subcortical Lesions and Parkinson's Disease*. *Cortex*, *36*(5), 601–622. http://doi.org/10.1016/S0010-9452(08)70541-0

Copland, D. A., Chenery, H. J., & Murdoch, B. E. (2001). Discourse Priming of Homophones in Individuals With Dominant Nonthalamic Subcortical Lesions, Cortical Lesions and Parkinsons Disease, *23*(4), 538–556. http://doi.org/10.1076/jcen.23.4.538.1233

Crinion, J., Turner, R., Grogan, A., Hanakawa, T., Noppeney, U., Devlin, J. T., et al. (2006). Language control in the bilingual brain. *Science*, *312*(5779), 1537–1540. http://doi.org/10.1126/science.1127761

Crosson, B., Benefield, H., Cato, M. A., Sadek, J. R., Moore, A. B., Wierenga, C. E., et al. (2003). Left and right basal ganglia and frontal activity during language generation: Contributions to lexical, semantic, and phonological processes. *Journal of the International Neuropsychological Society*, *9*(07), 1061–1077. http://doi.org/10.1017/S135561770397010X

de Diego Balaguer, R., Couette, M., Dolbeau, G., Dürr, A., Youssov, K., & Bachoud-Levi, A. C. (2008). Striatal degeneration impairs language learning: evidence from Huntington's disease. *Brain*, *131*(11), 2870–2881. http://doi.org/10.1093/brain/awn242

Duffy, S. A., Morris, R. K., & Rayner, K. (1988). Lexical ambiguity and fixation times in reading. *Journal of Memory and Language*, *27*(4), 429–446. http://doi.org/10.1016/0749-596X(88)90066-6

Duffy, S. A., Kambe, G., & Rayner, K. (2001). The effect of prior disambiguating context on the comprehension of ambiguous words: Evidence from eye movements.

Dunlosky, J., & Kane, M. J. (2007). The contributions of strategy use to working memory span: A comparison of strategy assessment methods. *The Quarterly Journal of Experimental Psychology*, *60*(9), 1227–1245. http://doi.org/10.1080/17470210600926075

Ehlen, F., Vonberg, I., Tiedt, H. O., Horn, A., Fromm, O., Kühn, A. A., & Klostermann, F. (2017). Thalamic deep brain stimulation decelerates automatic lexical activation. *Brain and Cognition*, *111*, 34–43. http://doi.org/10.1016/j.bandc.2016.10.001

Enard, W., Gehre, S., Hammerschmidt, K., & Cell, S. H. (2009). A humanized version of Foxp2 affects cortico-basal ganglia circuits in mice. *Elsevier*, *137*(5), 961–971. http://doi.org/10.1016/j.cell.2009.03.041

Fabbro, F., Clarici, A., & Bava, A. (1996). Effects of Left Basal Ganglia Lesions on Language Production. *Perceptual and Motor Skills*, *82*(3c), 1291–1298. http://doi.org/10.2466/pms.1996.82.3c.1291

Ferreira, F., Bailey, K. G. D., & Ferraro, V. (2016). Good-Enough Representations in Language Comprehension. *Current Directions in Psychological Science*, *11*(1), 11–15. http://doi.org/10.1111/1467-8721.00158

Fisher, S. E., & Scharff, C. (2009). FOXP2 as a molecular window into speech and language. *Trends in Genetics*, *25*(4), 166–177. http://doi.org/10.1016/j.tig.2009.03.002

Frank, M. J., Loughry, B., & O'Reilly, R. C. (2001). Interactions between frontal cortex and basal ganglia in working memory: A computational model. *Cognitive, Affective, & Behavioral Neuroscience*, *1*(2), 137–160. http://doi.org/10.3758/CABN.1.2.137

Frank, M. J., Seeberger, L. C., & O'Reilly, R. C. (2004). By Carrot or by Stick: Cognitive Reinforcement Learning in Parkinsonism. *Science*, *306*(5703), 1940–1943. http://doi.org/10.1126/science.1102941

Frank, M. J., Moustafa, A. A., Haughey, H. M., Curran, T., & Hutchison, K. E. (2007). Genetic triple dissociation reveals multiple roles for dopamine in reinforcement learning. Proceedings of the National Academy of Sciences, 104(41), 16311-16316.

Frank, M. J., & Hutchison, K. (2009). Genetic contributions to avoidance-based decisions: striatal D2 receptor polymorphisms. Neuroscience, 164(1), 131-140.

Friederici, A. D. (2006). What's in control of language? *Nature Neuroscience*, *9*(8), 991–992. http://doi.org/10.1038/nn0806-991

Friederici, A. D., Kotz, S. A., Werheid, K., & Hein, G. (2003). Syntactic comprehension in Parkinson's disease: investigating early automatic and late integrational processes using event-related brain potentials. ….

Frisch, S., Kotz, S. A., Cramon, Von, D. Y., & Friederici, A. D. (2003). Why the P600 is not just a P300: the role of the basal ganglia. *Clinical Neurophysiology*, *114*(2), 336–340. http://doi.org/10.1016/S1388-2457(02)00366-8

Fujita, E., Tanabe, Y., Shiota, A., Ueda, M., Suwa, K., Momoi, M. Y., & Momoi, T. (2008). Ultrasonic vocalization impairment of Foxp2 (R552H) knockin mice related to speech-language disorder and abnormality of Purkinje cells. *Proceedings of the National Academy of Sciences*, *105*(8), 3117–3122. http://doi.org/10.1073/pnas.0712298105

Gernsbacher, M. A., Varner, K. R., & Faust, M. E. (1990). Investigating differences in general comprehension skill. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *16*(3), 430–445. http://doi.org/10.1037/0278-7393.16.3.430

Gil Robles, S., Gatignol, P., Capelle, L., Mitchell, M.-C., & Duffau, H. (2005). The role of dominant striatum in language: a study using intraoperative electrical stimulations. *Journal of Neurology, Neurosurgery & Psychiatry*, *76*(7), 940–946. http://doi.org/10.1136/jnnp.2004.045948

Graybiel, A. M. (1995). Building action repertoires: memory and learning functions of the basal ganglia. *Current Opinion in Neurobiology*, *5*(6), 733–741. http://doi.org/10.1016/0959-4388(95)80100-6

Graybiel, A. M. (2000). The basal ganglia. *Current Biology*, *10*(14), R509–R511. http://doi.org/10.1016/S0960-9822(00)00593-5

Gregory Ashby, F., & Helie, S. (2011). A tutorial on computational cognitive neuroscience: Modeling the neurodynamics of cognition. *Journal of Mathematical Psychology*, *55*(4), 273–289. http://doi.org/10.1016/j.jmp.2011.04.003

Hazy, T. E., Frank, M. J., & O'Reilly, R. C. (2007). Towards an executive without a homunculus: computational models of the prefrontal cortex/basal ganglia system. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, *362*(1485), 1601–1613. http://doi.org/10.1098/rstb.2007.2055

Henson, R. (2005). What can functional neuroimaging tell the experimental psychologist?. The Quarterly Journal of Experimental Psychology Section A, 58(2), 193-233.

Hervais-Adelman, A. G., Moser-Mercer, B., & Golestani, N. (2011). Executive Control of Language in the Bilingual Brain: Integrating the Evidence from Neuroimaging to Neuropsychology. *Frontiers in Psychology*, *2*. http://doi.org/10.3389/fpsyg.2011.00234

Hervais-Adelman, A., Moser-Mercer, B., Michel, C. M., & Golestani, N. (2015). fMRI of Simultaneous Interpretation Reveals the Neural Basis of Extreme Language Control. *Cerebral Cortex*, *25*(12), 4727–4739. http://doi.org/10.1093/cercor/bhu158

Houk, J. C. (1997). Chapter 31 On the role of the cerebellum and basal ganglia in cognitive signal processing. *Progress in Brain Research*, *114*, 543–552. http://doi.org/10.1016/S0079-6123(08)63385-3

Järvikivi, J., Pyykkönen, P., & Niemi, J. (2009). Exploiting degrees of inflectional ambiguity: Stem form and the time course of morphological processing. Journal of Experimental Psychology: Learning, Memory, and Cognition, 35(1), 221.

Johns, B. T., Mewhort, D. J. K., & Jones, M. N. (2019). The Role of Negative Information in Distributional Semantic Learning. *Cognitive Science*, *43*(5), e12730. http://doi.org/10.1111/cogs.12730

Just, M. A., & Carpenter, P. A. (1992). A capacity theory of comprehension: Individual differences in working memory. *Psychological Review*, *99*(1), 122–149. http://doi.org/10.1037/0033-295X.99.1.122

Ketteler, D., Kastrau, F., Vohn, R., & Huber, W. (2008). The subcortical role of language processing. High level linguistic features such as ambiguity-resolution and the human brain; an fMRI study. *NeuroImage*, *39*(4), 2002–2009. http://doi.org/10.1016/j.neuroimage.2007.10.023

Kotseruba, I., & Tsotsos, J. K. (40). 40 years of cognitive architectures: core cognitive abilities and practical applications. Artificial Intelligence Review, 1-78.

Kotz, S. A., & Schmidt-Kassow, M. (2015). Basal ganglia contribution to rule expectancy and temporal predictability in speech. *Cortex*, *68*, 48–60. http://doi.org/10.1016/j.cortex.2015.02.021

Kotz, S. A., & Schwartze, M. (2010). Cortical speech processing unplugged: a timely subcortico-cortical framework. *Trends in Cognitive Sciences*, *14*(9), 392–399. http://doi.org/10.1016/j.tics.2010.06.005

Kotz, S. A., Frisch, S., & Cramon, Von, D. Y. (2003). Syntactic language processing: ERP lesion data on the role of the basal ganglia. *Journal of the International Neuropsychological Society*, *9*(07), 90. http://doi.org/10.1017/s1355617703970093

Kotz, S. A., Schwartze, M., & Schmidt-Kassow, M. (2009). Non-motor basal ganglia functions: A review and proposal for a model of sensory predictability in auditory language perception. *Cortex*, *45*(8), 982–990. http://doi.org/10.1016/j.cortex.2009.02.010

Kriete, T., Noelle, D. C., Cohen, J. D., & O'Reilly, R. C. (2013). Indirection and symbol-like processing in the prefrontal cortex and basal ganglia. *Proceedings of the National Academy of Sciences*, *110*(41), 16390–16395. http://doi.org/10.1073/pnas.1303547110

Kristensen, M., & Hansen, T. (2004). Statistical analyses of repeated measures in physiological research: a tutorial. Advances in physiology education, 28(1), 2-14.

Kuchinke, L., Hofmann, M. J., Jacobs, A. M., Frühholz, S., Tamm, S., & Herrmann, M. (2011). Human striatal activation during adjustment of the response criterion in visual word recognition. *NeuroImage*, *54*(3), 2412–2417. http://doi.org/10.1016/j.neuroimage.2010.08.062

Lai, C. S. L., Gerrelli, D., Monaco, A. P., Fisher, S. E., & Copp, A. J. (2003). FOXP2 expression during brain development coincides with adult sites of pathology in a severe speech and language disorder. *Brain*, *126*(11), 2455–2462. http://doi.org/10.1093/brain/awg247

Lieberman, P., Kako, E., Friedman, J., Tajchman, G., & Feldman, L. S. (1992). Speech production, syntax comprehension, and cognitive deficits in Parkinson's disease. *Brain and Language*. http://doi.org/10.1016/0093-934X(92)90127-Z

Liégeois, F., Baldeweg, T., Connelly, A., Gadian, D. G., Mishkin, M., & Vargha-Khadem, F. (2003). Language fMRI abnormalities associated with FOXP2 gene mutation. *Nature Neuroscience*, *6*(11), 1230–1237. http://doi.org/10.1038/nn1138

Lukatela, G., Frost, S. J., & Turvey, M. T. (1999). Identity priming in English is compromised by phonological ambiguity. Journal of Experimental Psychology: Human Perception and Performance, 25(3), 775.

MacDonald, M. C. (1994). Probabilistic constraints and syntactic ambiguity resolution. Language and cognitive processes, 9(2), 157-201.

Mason, R. A., & Just, M. A. (2007). Lexical ambiguity in sentence comprehension. *Brain Research*, *1146*, 115–127. http://doi.org/10.1016/j.brainres.2007.02.076

Marian, V., Blumenfeld, H. K., & Kaushanskaya, M. (2007). The Language Experience and Proficiency Questionnaire (LEAP-Q): Assessing language profiles in bilinguals and multilinguals. Journal of Speech, Language, and Hearing Research.

McNamara, D. S., & Scott, J. L. (2001). Working memory capacity and strategy use. *Memory & Cognition*, *29*(1), 10–17. http://doi.org/10.3758/BF03195736

Miller, E. K., & Cohen, J. D. (2001). An integrative theory of prefrontal cortex function. *Annual Review of Neuroscience*, *24*(1), 167–202. http://doi.org/10.1146/annurev.neuro.24.1.167

Mink, J. W. (1996). The basal ganglia: Focused selection and inhibition of competing motor programs. *Progress in Neurobiology*, *50*(4), 381–425. http://doi.org/10.1016/S0301-0082(96)00042-1

Moore, M., & Gordon, P. C. (2015). Reading ability and print exposure: Item response theory analysis of the author recognition test. Behavior research methods, 47(4), 1095-1109.

Newell, A. (1973). You can't play 20 questions with nature and win: Projective comments on the papers of this symposium.

Nixon, P., Lazarova, J., Hodinott-Hill, I., Gough, P., & Passingham, R. (2004). The inferior frontal gyrus and phonological processing: an investigation using rTMS. Journal of cognitive neuroscience, 16(2), 289-300.

Novick, J. M., Kan, I. P., Trueswell, J. C., & Thompson-Schill, S. L. (2009). A case for conflict across multiple domains: Memory and language impairments following damage to ventrolateral prefrontal cortex. Cognitive neuropsychology, 26(6), 527-567.

Osterhout, L., & Holcomb, P. J. (1992). Event-related brain potentials elicited by syntactic anomaly. *Journal of Memory and Language*, *31*(6), 785–806. http://doi.org/10.1016/0749-596X(92)90039-Z

Perfetti, C. (2007). Reading ability: Lexical quality to comprehension. Scientific studies of reading, 11(4), 357-383.

Prat, C. S., & Just, M. A. (2011). Exploring the neural dynamics underpinning individual differences in sentence comprehension. *Cerebral Cortex (New York, N.Y. : 1991)*, *21*(8), 1747–1760. http://doi.org/10.1093/cercor/bhq241

Reiner, A. (2009). You cannot have a vertebrate brain without a basal ganglia. In The basal ganglia IX (pp. 3-24). Springer, New York, NY.

Robertson, L. S. J., Corina, D. P., Ackerman, D., Guillemin, A., & Braun, A. R. (2004). Neural systems for sign language production: Mechanisms supporting lexical selection, phonological encoding, and articulation. *Human Brain Mapping*, *23*(3), 156–167. http://doi.org/10.1002/hbm.20054

Rodd, J. M., Davis, M. H., & Johnsrude, I. S. (2005). The Neural Mechanisms of Speech Comprehension: fMRI studies of Semantic Ambiguity. *Cerebral Cortex*, *15*(8), 1261–1269. http://doi.org/10.1093/cercor/bhi009

Rodd, J. M., Johnsrude, I. S., & Davis, M. H. (2010). The role of domain-general frontal systems in language comprehension: Evidence from dual-task interference and semantic ambiguity. *Brain and Language*, *115*(3), 182–188. http://doi.org/10.1016/j.bandl.2010.07.005

Rodd, J. M., Cutrin, B. L., Kirsch, H., Millar, A., & Davis, M. H. (2013). Long-term priming of the meanings of ambiguous words. Journal of Memory and Language, 68(2), 180-198.

Sambin, S., Teichmann, M., de Diego Balaguer, R., Giavazzi, M., Sportiche, D., Schlenker, P., & Bachoud Lévi, A. C. (2012). The role of the striatum in sentence processing: Disentangling syntax from working memory in Huntington's disease. *Neuropsychologia*, *50*(11), 2625–2635. http://doi.org/10.1016/j.neuropsychologia.2012.07.014

Schnur, T. T., Schwartz, M. F., Kimberg, D. Y., Hirshorn, E., Coslett, H. B., & Thompson-Schill, S. L. (2009). Localizing interference during naming: Convergent neuroimaging and neuropsychological evidence for the function of Broca's area. *Proceedings of the National Academy of Sciences*, *106*(1), 322–327. http://doi.org/10.1073/pnas.0805874106

Schultz, W., Dayan, P., & Montague, P. R. (1997). A Neural Substrate of Prediction and Reward. *Science*, *275*(5306), 1593–1599. http://doi.org/10.1126/science.275.5306.1593

Seidenberg, M. S., Tanenhaus, M. K., Leiman, J. M., & Bienkowski, M. (1982). Automatic access of the meanings of ambiguous words in context: Some limitations of knowledge-based processing. Cognitive psychology, 14(4), 489-537.

Senft, V., Stewart, T. C., Bekolay, T., Eliasmith, C., & Kröger, B. J. (2016). Reduction of dopamine in basal ganglia and its effects on syllable sequencing in speech: A computer simulation study. *Basal Ganglia*, *6*(1), 7–17. http://doi.org/10.1016/j.baga.2015.10.003

Seo, R., Stocco, A., & Prat, C. S. (2018). The bilingual language network: Differential involvement of anterior cingulate, basal ganglia and prefrontal cortex in preparation, monitoring, and execution. *NeuroImage*, *174*, 44–56. http://doi.org/10.1016/j.neuroimage.2018.02.010

Simmonds, A. J., Leech, R., Iverson, P., & Wise, R. J. S. (2014). The response of the anterior striatum during adult human vocal learning. *Journal of Neurophysiology*, *112*(4), 792–801. http://doi.org/10.1152/jn.00901.2013

Simmonds, A. J., Wise, R. J. S., Dhanjal, N. S., & Leech, R. (2011). A comparison of sensory-motor activity during speech in first and second languages. *Journal of Neurophysiology*, *106*(1), 470–478. http://doi.org/10.1152/jn.00343.2011

Stanovich, K. E., & West, R. F. (1989). Exposure to print and orthographic processing. Reading Research Quarterly, 402-433.

Stewart, A. J., Holler, J., & Kidd, E. (2007). Shallow processing of ambiguous pronouns: Evidence for delay. *The Quarterly Journal of Experimental Psychology*, *60*(12), 1680–1696. http://doi.org/10.1080/17470210601160807

Stewart, T. C., Bekolay, T., & Eliasmith, C. (2012). Learning to Select Actions with Spiking Neurons in the Basal Ganglia. *Frontiers in Neuroscience*, *6*. http://doi.org/10.3389/fnins.2012.00002

Stocco, A., & Anderson, J. R. (2008). Endogenous Control and Task Representation: An fMRI Study in Algebraic Problem-solving. *Journal of Cognitive Neuroscience*, *20*(7), 1300–1314. http://doi.org/10.1162/jocn.2008.20089

Stocco, A., & Prat, C. S. (2014). Bilingualism trains specific brain circuits involved in flexible rule selection and application. *Brain and Language*, *137*, 50–61. http://doi.org/10.1016/j.bandl.2014.07.005

Stocco, A., 2018. (2018). A Biologically Plausible Action Selection System for Cognitive Architectures: Implications of Basal Ganglia Anatomy for Learning and Decision‐Making Models. *Cognitive Science*, *42*, 457–490.

Stocco, A., Lebiere, C., & Anderson, J. R. (2010). Conditional routing of information to the cortex: a model of the basal ganglia's role in cognitive coordination. *Psychological Review*, *117*(2), 541–574. http://doi.org/10.1037/a0019077

Stocco, A., Lebiere, C., O'Reilly, R. C., & Anderson, J. R. (2012). Distinct contributions of the caudate nucleus, rostral prefrontal cortex, and parietal cortex to the execution of instructed tasks. *Cognitive, Affective, & Behavioral Neuroscience*, *12*(4), 611–628. http://doi.org/10.3758/s13415-012-0117-7

Stocco, A., Yamasaki, B., Natalenko, R., & Prat, C. S. (2014). Bilingual brain training: A neurobiological framework of how bilingual experience improves executive function. *International Journal of Bilingualism*, *18*(1), 67–92. http://doi.org/10.1177/1367006912456617

Stocco, A., Murray, N. L., Yamasaki, B. L., Renno, T. J., Nguyen, J., & Prat, C. S. (2017). Individual differences in the Simon effect are underpinned by differences in the competitive dynamics in the basal ganglia: An experimental verification and a computational model. Cognition, 164, 31-45.

Sutton, R. S., & Barto, A. G. (1987, July). A temporal-difference model of classical conditioning. In Proceedings of the ninth annual conference of the cognitive science society (pp. 355-378).

Sutton, R. S., & Barto, A. G. (1998). Introduction to reinforcement learning. Vol. 135.

Swinney, D. A. (1979). Lexical access during sentence comprehension:(Re) consideration of context effects. Journal of verbal learning and verbal behavior, 18(6), 645-659.

Tabossi, P., & Zardon, F. (1993). Processing ambiguous words in context. Journal of memory and language, 32(3), 359-372.

Takahashi, K., Liu, F. C., Hirokawa, K., & Takahashi, H. (2003). Expression of Foxp2, a gene involved in speech and language, in the developing and adult striatum. *Journal of Neuroscience Research*, *73*(1), 61–72. http://doi.org/10.1002/jnr.10638

Tan, L. H., Chen, L., Yip, V., Chan, A. H. D., Yang, J., Gao, J.-H., & Siok, W. T. (2011). Activity levels in the left hemisphere caudate–fusiform circuit predict how well a second language will be learned. *Proceedings of the National Academy of Sciences*, *108*(6), 2540–2544. http://doi.org/10.1073/pnas.0909623108

Tettamanti, M., Moro, A., Messa, C., Moresco, R. M., Rizzo, G., Carpinelli, A., et al. (2005). Basal ganglia and language: phonology modulates dopaminergic release. *NeuroReport*, *16*(4), 397.

Twilley, L. C., Dixon, P., Taylor, D., & Clark, K. (1994). University of Alberta norms of relative meaning frequency for 566 homographs. *Memory & Cognition*, *22*(1), 111–126. http://doi.org/10.3758/BF03202766

Ullman, M. T. (2001). A neurocognitive perspective on language: The declarative/procedural model. *Nature Reviews Neuroscience*, *2*(10), 717–726. http://doi.org/10.1038/35094573

Ullman, M. T. (2004). Contributions of memory circuits to language: the declarative/procedural model. *Cognition*, *92*(1-2), 231–270. http://doi.org/10.1016/j.cognition.2003.10.008

Ullman, M. T. (2016). The Declarative/Procedural Model: A Neurobiological Model of Language Learning, Knowledge, and Use. In *Neurobiology of Language* (pp. 953–968). Academic Press. http://doi.org/10.1016/B978-0-12-407794-2.00076-6

Ullman, M. T., & GOPNIK, M. (1999). Inflectional morphology in a family with inherited specific language impairment. *Applied Psycholinguistics*, *20*(1), 51–117.

Ullman, M. T., Corkin, S., Coppola, M., Hickok, G., Growdon, J. H., Koroshetz, W. J., & Pinker, S. (2008). A Neural Dissociation within Language: Evidence that the Mental Dictionary Is Part of Declarative Memory, and that Grammatical Rules Are Processed by the Procedural System. *Dx.Doi.org*, *9*(2), 266–276. http://doi.org/10.1162/jocn.1997.9.2.266

van der Lely, H. K. J., & Christian, V. (2000). Lexical word formation in children with grammatical SLI: a grammar-specific versus an input-processing deficit? *Cognition*, *75*(1), 33–63. http://doi.org/10.1016/S0010-0277(99)00079-7

Van Hasselt, H., Guez, A., & Silver, D. (2016, March). Deep reinforcement learning with double q-learning. In Thirtieth AAAI Conference on Artificial Intelligence.

Vargha-Khadem, F., Gadian, D. G., Copp, A., & Mishkin, M. (2005). FOXP2 and the neuroanatomy of speech and language. *Nature Reviews Neuroscience*, *6*(2), 131–138. http://doi.org/10.1038/nrn1605

Vargha-Khadem, F., Watkins, K. E., Price, C. J., Ashburner, J., Alcock, K. J., Connelly, A., et al. (1998). Neural basis of an inherited speech and language disorder. *Proceedings of the National Academy of Sciences*, *95*(21), 12695–12700. http://doi.org/10.1073/pnas.95.21.12695

Vitello, S., & Rodd, J. M. (2015). Resolving Semantic Ambiguities in Sentences: Cognitive Processes and Brain Mechanisms. *Language and Linguistics Compass*, *9*(10), 391–405. http://doi.org/10.1111/lnc3.12160

Vitello, S., Warren, J. E., Devlin, J. T., & Rodd, J. M. (2014). Roles of frontal and temporal regions in reinterpreting semantically ambiguous sentences. *Frontiers in Human Neuroscience*, *8*(416), 1124. http://doi.org/10.3389/fnhum.2014.00530

Vonberg, I., Ehlen, F., Fromm, O., Kühn, A. A., & Klostermann, F. (2016). Deep Brain Stimulation of the Subthalamic Nucleus Improves Lexical Switching in Parkinson's Disease Patients. *PLoS ONE*, *11*(8), e0161404. http://doi.org/10.1371/journal.pone.0161404

Wallesch, C.-W., Henriksen, L., Kornhuber, H. H., & Paulson, O. B. (1985). Observations on regional cerebral blood flow in cortical and subcortical structures during language production in normal man. *Brain and Language*, *25*(2), 224–233. http://doi.org/10.1016/0093-934X(85)90082-3

Watkins, K. E., Vargha-Khadem, F., Ashburner, J., Passingham, R. E., Connelly, A., Friston, K. J., et al. (2002). MRI analysis of an inherited speech and language disorder: structural brain abnormalities. *Brain*, *125*(3), 465–478. http://doi.org/10.1093/brain/awf057

Yin, H. H., & Knowlton, B. J. (2006). The role of the basal ganglia in habit formation. *Nature Reviews Neuroscience*, *7*(6), 464–476. http://doi.org/10.1038/nrn1919