**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**

Andrea Stocco^{1,2,3,4*} Nicole L. Murray², Brianna L. Yamasaki^{1,2},
Taylor J. Renno¹, Jimmy Nguyen¹, and Chantel S. Prat^{1,2,3,4}

¹ Department of Psychology, University of Washington, Seattle, WA 98195

² Institute for Learning and Brain Sciences (I-LABS), University of Washington, Seattle, WA
98195

³ Center for Sensorimotor Neural Engineering, University of Washington, Seattle, WA 98195

⁴ University of Washington Institute of Neuroengineering (UWIN), Seattle, WA 98195

* Corresponding Author: Campus Box 351575, University of Washington, Seattle, WA 98195,
Email: stocco@uw.edu

Email Addresses: AS: stocco@uw.edu; NLM: nlmurray44@hotmail.com; BLY: bly4@uw.edu;
TJR: tayoren@uw.edu; JN: jimmynguyen431@gmail.com; CSP: csprat@uw.edu

Abstract

Cognitive control is thought to be made possible by the activity of the prefrontal cortex, which selectively uses task-specific representations to bias the selection of task-appropriate responses over more automated, but inappropriate, ones. Recent models have suggested, however, that prefrontal representations are in turn controlled by the basal ganglia. In particular, neurophysiological considerations suggest that the basal ganglia's indirect pathway plays a pivotal role in preventing irrelevant information from being incorporated into a task, thus reducing response interference due to the processing of inappropriate stimuli dimensions. Here, we test this hypothesis by showing that individual differences in a non-verbal cognitive control task (the Simon task) are correlated with performance on a decision-making task (the Probabilistic Stimulus Selection task) that tracks the contribution of the indirect pathway. Specifically, the higher the effect of the indirect pathway, the smaller was the behavioral costs associated with suppressing interference in incongruent trials. Additionally, it was found that this correlation was driven by individual differences in incongruent trials only (with little effect on congruent ones) and specific to the indirect pathway (with almost no correlation with the effect of the direct pathways). Finally, it is shown that this pattern of results is precisely what is predicted when competitive dynamics of the basal ganglia are added to the selective attention component of a simple model of the Simon task, thus showing that our experimental results can be fully explained by our initial hypothesis.

Keywords: Cognitive control; Selective attention; Prefrontal cortex; Basal ganglia; Computational modeling

1. Introduction

Cognitive control depends on the capacity to overcome prepotent behavioral responses that interfere with those required by internal goals. In many cases, the successful resolution of this interference depends on *selective attention*, that is, the capacity to ignore certain features of a stimulus and instead focus on other characteristics, selected on the basis of internal task goals. However, the exact nature of the neural mechanism for resolving interference is still debated. In fact, multiple mechanisms might be recruited for cognitive control, depending on the nature and demands of the task at hand. In this paper, we suggest that the resolution of interference in a task that requires cognitive control relies on the activity of the basal ganglia, a subcortical circuit believed to be responsible for selecting which sensory information is ultimately transmitted to the prefrontal cortex (PFC). We then present an empirical test of this hypothesis and a computational model that accounts for the data presented herein.

The remainder of this paper is structured in four parts. First, we provide a brief overview of the problem of cognitive control in response selection interference. Second, we introduce the basal ganglia circuit, its role in response selection, and a behavioral task that measures the competitive basal ganglia dynamics to resolve conflict. Third, we present an experiment that demonstrates the existence of a significant (and hitherto unsuspected) correlation between performance on a cognitive control task (the Simon task: Craft & Simon, 1970; Simon, 1990) and a task that measures competitive basal ganglia dynamics (the Probabilistic Stimulus Selection task: Frank, Seeberger, & O'Reilly, 2004). Lastly, we present a computational model that provides an explicit and mechanistic account of interference in the Simon task in terms of basal ganglia dynamics, which accounts for our results.

1.1 Cognitive Control

Cognitive control is typically studied through tasks that require managing interference between competing responses. For example, in the Stroop task, participants are asked to say out loud the name of the color a word is presented in while ignoring the word itself. When the word is the name of a different color (e.g., RED printed in blue), interference arises because of the tendency to name the word. This interference is manifested in longer reaction times to such incongruent trials, as opposed to congruent ones where no interference exists (e.g., RED printed in red). Another example is the Simon task (Craft & Simon, 1970; Lu & Proctor, 1995; Proctor & Lu, 1999; Simon, 1990), in which participants are asked to respond with their left and right hand to specific visual features (e.g., shape or color) of a stimulus that appears on a screen. For example, they might be asked to respond with their left hand when the stimulus is a square, and with their right hand if the stimulus is a circle (Figure 1). Interference occurs when a stimulus is presented on the side of the screen that is contralateral to the desired response. As a result, these “incongruent” trials (Figure 1C-D) are less accurate and take longer than “congruent” trials in which the stimulus appears on the side of the desired response (Figure 1A-B). This extra time is supposed to reflect the additional cost necessary to resolve conflict generated by the activation of two competing responses (one for the shape, one for the position).

Figure 1: Overview of the Simon task. In the case illustrated here, participants are instructed to respond with their left hand if the stimulus is a square, and with their right hand if it is a circle. Additionally, the stimulus can appear on the left or the right side of the screen. During congruent trials (A and B), the stimuli appear on the same side as their associated response. During incongruent trials (C and D), on the other hand, the stimuli appear on the opposite side, thus creating a competition between the correct response associated with the shape of the stimulus and the incorrect response associated with the location of the stimulus. Incongruent trials are typically associated with more errors and longer response times (Figure 4).

The precise source of interference in these tasks has been much debated and might vary across different paradigms (Liu, Banich, Jacobson, & Tanabe, 2004; Nee, Wager, & Jonides, 2007; van Veen & Carter, 2005). In the specific case of the Simon task, interference likely occurs early on, at the moment in which the relevant and irrelevant features of the stimuli are being processed, as evidenced from both a review of the behavioral data (Lu & Proctor, 1995) and from the onset of brain oscillations in the fronto-parietal network that reliably indicate individual differences in working memory encoding (Gulbinaite, van Rijn, & Cohen, 2014).

While different authors might disagree on the source of interference for specific tasks, they tend to agree that, at the neural level, interference is resolved through mechanisms underpinned by the prefrontal cortex (Miller, 2000; Miller & Cohen, 2001). More specifically, by holding a representation of the intended goal (for instance, paying attention to the shape of the stimulus, rather than its position), PFC exerts a top-down influence that ultimately counters the prepotency of unwanted responses. The role of PFC in exerting this form of control has been verified in numerous imaging studies (Cole, Laurent, & Stocco, 2013; Cole & Schneider, 2007; Kane & Engle, 2002; Koechlin, Ody, & Kouneiher, 2003; MacDonald, Cohen, Stenger, & Carter, 2000; Miller, 2000; Ridderinkhof, van den Wildenberg, Segalowitz, & Carter, 2004), proposed for both proactive and reactive control (Braver, 2012; Braver, Paxton, Locke, & Barch, 2009) and implemented in numerous neurocomputational models (Botvinick & Plaut, 2004; Cohen, Braver, & O'Reilly, 1996; Cohen, Dunbar, & McClelland, 1990; Herd, Banich, & O'Reilly, 2006). Within neural network models, this mechanism is typically implemented by adding specific “task” or “goal” representational units at the top of the network hierarchy. The activation spreading from these units ultimately provides the necessary boost to overcome the interference caused by irrelevant stimuli that are encoded by units whose synaptic weights are normally much stronger. For example, in the early influential model of the Stroop task by Cohen et al. (1990), different sets of units encode the stimulus’ word and the stimulus’ color, and both are connected to the output layer of the network, whose units represent color names. Stroop interference originates because the connectivity between input *word* units and output color names is much stronger than that between input *color* units and output color names. The model can successfully overcome interference, however, when additional input units are added that encode the task goal (that is, whether to name the color or read the word).

For such a mechanism to function properly, PFC must be able to maintain active representations of the relevant information, while at the same time discarding irrelevant information that would lead to a top-down activation of the irrelevant or incorrect responses. Indeed, single-cell recordings show that PFC neurons exhibit precisely this type of selectivity. For example, PFC neurons were found to encode only the relevant information of the stimulus (the location of a target) and not other aspects (i.e., the location of distractors; Rainer, Asaad, & Miller, 1998). Conversely, accidental encoding of irrelevant information in PFC is associated with poor performance in cognitive control tasks. For instance, compared to individuals with lower cognitive capacity (as indexed by working memory measures), individuals with higher cognitive capacity exhibit smaller event-related potentials when asked to remember the location of a visual target stimulus that is surrounded by distractors, suggesting that better cognitive abilities are associated with a greater capacity to ignore irrelevant information (Vogel, McCollough, & Machizawa, 2005). In turn, this suggests a tight relationship between PFC, working memory, and selective attention, whereby working memory depends on the successful encoding of information in PFC, which ultimately relies on selective attention to optimally allocate PFC resources to the relevant characteristics of the stimuli (Awh, Vogel, & Oh, 2006; Engle, Kane, & Tuholski, 1999; Kane & Engle, 2002, 2003; McNab & Klingberg, 2008).

1.2 Basal Ganglia Competitive Dynamics of Access to PFC

The theory that interference is resolved through top-down representations provides an elegant solution to the problem of cognitive control and has been largely confirmed, but it poses a second problem: How are these representations selected? In the past two decades, much evidence has accumulated suggesting that the access to PFC is largely regulated by the basal ganglia, a set of interconnected subcortical nuclei (Figure 2) that receive converging projections

from the entire cortex, but project selectively to the frontal lobes, thus operating as a “funnel” through which cortical signals must compete for access to PFC (Albin, Young, & Penney, 1989; Alexander, DeLong, & Strick, 1986; DeLong, 1990; Kemp & Powell, 1971). Because of their peculiar anatomy, it has been repeatedly suggested that one of the functions of the basal ganglia is to select among competing signals that need to be relayed (through the thalamus) to PFC (Gurney, Prescott, & Redgrave, 2001; Redgrave, Prescott, & Gurney, 1999; Stephenson-Jones, Samuelsson, Ericsson, Robertson, & Grillner, 2011). Although this function was originally proposed to describe the selection of motor programs (Redgrave et al., 1999), more recent neurocomputational models have suggested that the same mechanism can be generalized to cognitive processes (Frank, Samanta, Moustafa, & Sherman, 2007; Houk et al., 2007; Humphries, Stewart, & Gurney, 2006; O’Reilly & Frank, 2006; Stewart, Choo, & Eliasmith, 2010; Stocco & Lebiere, 2014; Stocco, Lebiere, & Anderson, 2010). For example, O’Reilly and Frank (2006) have shown that, in a neural network model of the interactions between basal ganglia and PFC, the basal ganglia learn how to adaptively gate sensory signals to PFC for future elaboration. In other words, the basal ganglia learn to adaptively and strategically use PFC as a working memory buffer and to control which stimuli are worth memorizing (“gating”) and which, instead, can be discarded.

The most commonly held view is that the selection of sensory signals is mediated by the competition between at least two pathways within the circuit (Albin et al., 1989; DeLong, 1990). These two pathways have been traditionally indicated as the *direct* and the *indirect* pathway, and have excitatory and inhibitory effects, respectively, on thalamic neurons projecting to the PFC (Figure 2). Specifically, the direct pathway results in an *increase* of thalamic inputs to the cortex, while the indirect pathway results in a net *decrease* of these inputs.

Crucially, the opposition between the two pathways is mediated by dopamine-driven reinforcement-learning mechanisms (Niv, Duff, & Dayan, 2005; Schultz, 1998; Schultz, Apicella, & Ljungberg, 1993; Schultz, Dayan, & Montague, 1997), so that the basal ganglia effectively learns to optimize the selection of the most relevant signals to send to PFC based on their expected utility (O'Reilly & Frank, 2006; Stocco et al., 2010). Some models have also outlined a connection between the activity of the two pathways and specific cognitive abilities. Specifically, the opposite effects of the direct and indirect pathways on the thalamic projections has been likened to controlling a “gate” (Frank, Loughry, & O'Reilly, 2001; O'Reilly & Frank, 2006) or a “router” (Stocco, 2012; Stocco & Lebiere, 2014; Stocco et al., 2010) to PFC, thus determining which stimuli will be maintained, and which will be updated, in working memory.

Because the indirect pathway has an inhibitory effect on thalamic projections, its role becomes more prominent as cognitive performance becomes more dependent on discarding irrelevant task features. For example, increased expression of dopamine D2 receptors, which are specific to the indirect pathway (Gerfen et al., 1990), is associated with better performance in an *N*-back task (Zhang et al., 2007). Because this task requires participants to maintain a moving window of the last *N* items in a stream of stimuli that are presented, it places severe demands on controlling which information is added to the contents of working memory. Similarly, McNab and Klingberg (2008) found that basal ganglia activity, most likely associated with the indirect pathway (because it is localized in the globus pallidus: Figure 2), correlates with better performance in a visual working memory task in which participants are required to memorize the location of certain stimuli (e.g., red dots) while ignoring distractors (e.g., yellow dots). Conversely, reduced expression of D2 receptors has been associated with reduced visuospatial skills (Berman & Noble, 1995) and impaired performance on the Wisconsin Card Sorting Task

(Han et al., 2008). Finally, at least one study reported that administration of cabergoline, a D2-agonist, resulted in improved performance in response inhibition and executive function tasks (Messer, 2011).

In summary, a vast amount of research suggests that the basal ganglia play a fundamental role in selecting which sensory information is ultimately encoded in PFC, and therefore, which dimension of a stimulus will be used to resolve conflict between competing responses. Furthermore, neurophysiological evidence and computational models jointly suggest that the key function of filtering out unwanted information is specific to one of two pathways within the basal ganglia circuit, the so-called indirect pathway. For this reason, we expect that individual differences in cognitive control should be related to individual differences in the effects of the indirect pathway on PFC.

Figure 2: An overview of the basal ganglia circuit and the direct and indirect pathways within the circuit. For clarity purposes, the figure omits existing backward projections within the circuit (e.g., from thalamus to striatum), as well as the hyperdirect pathway.

1.3 The Probabilistic Stimulus Selection (PSS) Task

Because the two pathways are interleaved throughout the basal ganglia, their competitive effects cannot be dissociated easily. However, their relative contribution can be indirectly measured through the Probabilistic Stimulus Selection (PSS) task (Frank, Seeberger, & O'Reilly, 2004). The PSS task is an iterative, implicit decision-making paradigm in which participants repeatedly choose from pairs of non-verbalizable stimuli, each of which has a different probability of yielding a reward. In the PSS task, participants are initially trained to select the most rewarding stimulus out of three different pairs (Figure 3A-B). Note that the correct response can be made by either learning to choose the most rewarding stimulus, or by learning to avoid the least rewarding one. To distinguish between these two strategies, participants are then tested over the remaining combinations of stimuli (Figure 3C-D), so that their sensitivity for learning to choose high-reward probability stimuli ("Choose" accuracy) and their sensitivity for learning to avoid low-reward probability stimuli ("Avoid" accuracy) can be measured independently. Multiple patient, pharmacological, and genetic studies (Frank et al. 2004; Frank & Hutchinson, 2009; Frank et al., 2007) have shown that Choose accuracy reflects the contribution of the direct pathway, while Avoid accuracy reflects that of the indirect pathway.

Figure 3. Overview of the Probabilistic Stimulus Selection task. During the training phase, participants repeatedly select one stimulus from three possible pairs. (A) During the training phase, choices always result in a probabilistic feedback, which is conveyed immediately after each choice. The probability of positive feedback (“Correct!”) depends on the stimulus being chosen. (B) Stimuli are always presented in fixed pairs (dotted lines); within each pair, one stimulus is more likely than the other to yield a positive outcome. (C) During the test phase, participants perform more choices but do not receive any more feedback. (D) In addition, during the test phase, the original pairs are separated, and participants are now asked to decide between pairs generated by combining either the most rewarding stimulus (green lines) or the least rewarding stimulus (red lines) against each of the remaining stimuli.

For example, Frank et al. (2007) showed that better “Choose” accuracy was associated with polymorphisms in the DARP32 gene, known to affect the expression of D1 dopamine receptors in the direct pathway, while “Avoid” accuracy was associated with polymorphisms in the C957T gene, known to affect the expression of D2 dopamine receptors in the indirect pathway (Frank et al., 2007; Frank et al., 2009). Additional evidence comes from patients with

Parkinson's Disease, a pathology that results from the death of cells releasing dopamine in the striatum. Because dopamine excites the direct pathway but inhibits the indirect pathway, Parkinson's Disease results in abnormal activation of the indirect pathway at the expense of the direct pathway. which is paralleled by a corresponding increase in Avoid accuracy and decrease in Choose accuracy (Frank et al., 2004). Finally, when Parkinson's patients were tested after taking dopamine-promoting medication (and thus increasing the activity of the direct pathway at the expense of the indirect one), their Choose accuracy increased while their Avoid accuracy plummeted.

1.5 Summary and Predictions

In summary, the existing literature suggests that successful cognitive control depends on the proper allocation of attention to the relevant dimensions of a stimulus and that this process, in turn, depends at least partially on the activity of the basal ganglia, and of the indirect pathway in particular. If this assumption is correct, individual differences in cognitive control should be correlated with individual differences in the activity of the indirect pathway. To test this prediction, we conducted an experiment to correlate individual performance in the PSS task (whose accuracies are affected by the relative strengths of the direct and indirect pathway) and the size of the incongruency effect in the Simon task (which reflects a person's capacity for cognitive control). To control for confounding factors, independent measures of working memory capacity were also collected and partialled out. Our main prediction was that, at the group level, Avoid accuracy in the PSS task should be negatively correlated with the incongruency effect in the Simon task. Furthermore, we expect that this correlation is selectively due to the capacity to filter out the irrelevant stimulus dimension (that is, spatial position) during the incongruent trials. Thus, we expect that a similar correlation exists between mean response

times to incongruent trials and Avoid accuracy in the PSS, but not between Avoid accuracy and response times for congruent trials. Finally, because we expect that the act of processing both stimulus shape and position are highly automatic and unlikely to change over the course of a single experimental session, we do not expect to find any significant correlation between Choose accuracy and response times in the Simon task.

2. Materials and Methods

2.1 Participants

Fifty-eight healthy individuals were recruited for the experiment (age = 18-34 years, 44 females). Data from 8 participants (5 female) were not analyzed due to an inability to attain learning criteria required during the PSS task learning phase. All participants were recruited from the student population of the University of Washington campus and the surrounding Seattle area and received monetary compensation in exchange for their time. All participants provided written informed consent in accordance with the ethical guidelines established by the IRB prior to the start of the experiment.

2.2 Procedures

The PSS task, the Simon task, and one test of complex working memory span (the Operation Span task: Unsworth, Heitz, Schrock, & Engle, 2005) were administered to participants as part of a larger cognitive battery assessing cognitive capabilities in participants who went on to participate in a training experiment not reported herein. With the exception of the three tasks described herein, no other experimental task was selected for this specific analysis, and no dataset was discarded after having been considered for this study. Twenty-eight participants completed the Simon task prior to the PSS task while 22 participants completed the PSS task prior to the Simon task. The Simon and PSS tasks were executed in different sessions

on separate days. Both tasks were performed on a computer, in front of a 21" LCD screen, using a standard keyboard to respond. Stimuli presentation and response collection were controlled through the E-prime software (Psychological Software Tools, Pittsburgh, PA).

2.2.1 The Simon Task

Behavioral measures of cognitive control were collected through the Simon task (Craft & Simon, 1970; Simon, 1990). Before the beginning of the task, participants were asked to position their left index finger on the "Q" key of the keyboard, and their right index finger on the "7" key of the keyboard's numerical pad. During the task, participants were presented with one of two shapes, either a black square or a black circle, on a white background. Participants were instructed to respond to one shape (e.g., squares) with their right hand, and to the other shape (e.g., circles) with their left hand; the associations between shapes and response keys was randomized across participants. Each trial was introduced by an 800 ms fixation, followed by a 250 ms gap, followed by the stimuli. The stimuli remained on the screen for either 3000 ms or until a response was recorded. A new trial began immediately afterward.

The shapes were presented in randomized trials that were either congruent (e.g., a stimulus associated with a left response and presented on the left half of the screen) or incongruent (e.g., a stimulus associated with a left response and presented on the right half of the screen). The congruent to incongruent trial ratio was 3:1, with 75% of the trials being congruent and 25% of trials being incongruent. Both congruent and incongruent trials consisted of an equal number of circles and squares. A total of 200 trials were presented, divided into 4 blocks of 50 trials each, and presented in random order.

2.2.2 The Probabilistic Stimulus Selection (PSS) Task

Behavioral measures of basal ganglia function were collected using the PSS task. The computerized version of the task was the same used in Frank, Seeberger, and O'Reilly's (2004) seminal paper. The task consists of two consecutive phases, a training phase (Figure 3A/B) and a test phase (Figure 3C/D). In both phases, participants performed multiple decision-making trials in which they were asked to choose one of two Japanese Hiragana characters, placed at the left and right side of the screen. The use of Hiragana characters was intended to discourage participants from memorizing the stimuli through verbal rehearsal strategies and encourage them to make their decisions intuitively (Frank, Seeberger, & O'Reilly, 2004). None of our participants were familiar with the Japanese writing system before the experiment. Participants indicated their response by pressing the keys "1" (for the left character) and "0" (for the right character) on a standard keyboard. Each choice resulted in a binary feedback (either "Correct!" or "Incorrect"), whose probability depended on the stimulus chosen. A total of six stimuli were sampled, each of which was associated with a unique probability of success (Figure 3B). The association between each character and its success probability was randomized for each participant but remained consistent across phases.

During the training phase, the six stimuli were presented in three fixed pairs, which were presented 20 times each for a total of 60 trials. To ensure that participants could discriminate the relative success probability for each stimulus, the training phase was repeated until each subject's accuracy had reached a predetermined criterion, measured as the percentage of times the most rewarding stimulus in each pair was chosen. As in the original study (Frank et al. 2004), the criterion was different for each pair (65%, 60%, and 50%, respectively, for the top, middle, and bottom pairs in Figure 3B) to account for the different probabilities of each stimulus being

correct. After a maximum of six repetitions of the training phase, participants moved on to the test phase. As in previous experiments with this task (Frank, Moustafa, Haughey, Curran, & Hutchison, 2007; Frank et al., 2004), participants who did not reach the learning criterion by the end of the sixth phase were excluded from the analysis.

Critically, during the training phase, participants can learn to decide by either *choosing* the most rewarding stimulus or by *avoiding* the least rewarding one. To distinguish between these two strategies, the test phase presents each of the 15 possible combinations of stimuli four times, for a total of 60 trials (Figure 3D). The decisions made during the testing phase yield two distinct measures of performance: the accuracy in choosing the most rewarding stimulus against all the others (“Choose” accuracy, i.e., the proportion of choices in which stimulus *A* was preferred over *C*, *D*, *E* and *F*) and the accuracy in avoiding the least rewarding stimulus when it was paired with all the other (“Avoid” accuracy, or the proportion of choices in which stimuli *C*, *D*, *E*, and *F* were preferred over *B*).

2.2.3 The Operation Span Task

Individual differences in working memory capacity were assessed using the automated version of the Operation Span task (Unsworth et al., 2005), a non-verbal test of complex working memory span. In the Operation Span task, participants performed a series of trials that required memorizing a sequence of letters presented in alternation to the evaluation of arithmetic expressions. In each trial, an arithmetic expression was first presented, such as “ $(2 \times 5) - 9 =$ ”, followed by a probe number (e.g., “2”). Participants had to indicate whether the probe was the result of the arithmetic expression by clicking on two “yes” or “no” buttons on the screen. After their response, one letter was briefly presented on the screen. A new arithmetic expression and a new letter were then presented. After the presentation of the last letter, participants were required to indicate all letters that had been presenting since the beginning of the trial by selecting them in the order that they appeared

from an array of 4 x 3 letters. The number of letters to be memorized varied pseudo-randomly across trials, up to a maximum number of seven (Unsworth et al., 2005). Working memory capacity was measured in terms of the *span score* (Unsworth et al., 2005), which ranges from a minimum of 0 to a maximum of 75, the latter indicating that all of the trials were answered correctly.

3. Experimental Results

3.1 General Results in the PSS Task

A *t*-test of Choose and Avoid conditions revealed no group-level mean differences in accuracy [$t(49) = 1.15$, $p = 0.25$; Figure 4A]. Additionally, consistent with previous studies (Frank & Hutchinson, 2009), Choose and Avoid mean accuracies were not significantly related with one another [$r(50) = 0.02$, $p = 0.88$].

3.2 General Results in the Simon Task

As expected, the mean response time for incongruent trials ($M = 489$ ms, $SD = 88$ ms) was significantly longer than for congruent trials ($M = 421$ ms, $SD = 73$ ms, paired $t(49) = 10.91$, $p < 0.0001$; Figure 4C). Because accuracy data is not normally distributed, the data was arcsine-root transformed before being analyzed. The two conditions also differed significantly in terms of accuracy ($t(49) = 6.32$, $p < 0.0001$), with incongruent trials being significantly less accurate ($M = 0.88$, $SD = 0.03$) than congruent trials ($M = 0.98$, $SD = 0.10$; Figure 4B).

Summary of Behavioral Measures

Figure 4: Group results from the PSS Task (left) and the Simon task (center and right). Bars represent mean values; error bars represent standard errors of the mean. * $p < 0.0001$.

3.3. General Results in the Operation Span Task

The mean working memory capacity score on the Operation Span task was 43.34 ($SD = 17.75$), which is in line with the results of previous studies (Unsworth & Engle, 2007; Unsworth et al., 2005). Individual differences in Operation Span scores were not correlated with any other behavioral measure of the other two tasks ($r^2 < 0.01$, $p > 0.33$).

3.4 Correlations Between Avoid and Choose Accuracies and Simon Task Measures

Two-tailed correlational analyses were conducted relating group level data with both Choose and Avoid accuracy conditions with congruent and incongruent reaction times and congruency effect (incongruent - congruent) on the Simon task. As expected, a significant

negative correlation was found between Avoid accuracy and response times for incongruent trials [$r(50) = -0.40$, $p = 2.84$, $p = 0.004$]. Avoid accuracy was also correlated with the size of the Simon effect [that is, the difference between incongruent and congruent response times: $r(50) = -0.45$, $p = 0.001$]. Consistent with our expectations, no significant correlation was found between Avoid accuracy and congruent response times [$r(50) = -0.22$, $p = 0.12$]. No significant correlations were found between Choose accuracy and response times for either congruent or incongruent trials, or their difference [$r(50) < 0.10$, $p > 0.40$]. The results of these correlations are reported in Figure 5.

To exclude any confound due to working memory capacity, the same correlations were calculated after partialling out the Operation Span scores. None of the correlations were affected by this procedure. Specifically, after including the span scores, the partial correlation between Avoid accuracy and incongruent trials was still $r(48) = -0.40$, $p = 0.004$; the partial correlation between Avoid accuracy and Simon effect was $r(48) = -0.44$, $p = 0.002$; and the partial correlation between Avoid accuracy and congruent trials was $r(48) = -0.23$, $p = 0.11$. Similarly, all the correlations between Choose accuracy and Simon task response times remained non-significant ($p > 0.45$). This pattern suggests that the amount of variability in Simon task performance that is explained by working memory is much smaller than the variability explained by PSS accuracies (and, by extension, by the subcortical contributions of the basal ganglia).

Correlations between PSS Task and Simon Task

Figure 5: Correlations between Choose and Avoid accuracies in the PSS task (left and right columns) and response times in the Simon task (rows). Dotted lines represent regression lines, shaded areas represent 95% confidence intervals of the estimates. Red lines and shaded areas represent significant correlations; grey lines and shaded areas represent non-significant ones.

Finally, both the correlation between Avoid accuracy and incongruent response times and the correlation between Avoid accuracy and the interference effect were significantly higher than their counterparts with Choose accuracy (Fisher's $Z = 2.44$, $p = 0.01$ and $Z = 2.25$, $p = 0.02$).

4. Computational Model

As hypothesized, our results have shown that the size of the Simon effect, believed to reflect the additional cost of cognitive control, is related to the Avoid accuracy in the PSS task, a measure that is related to the activity of the indirect pathway in the basal ganglia. Specifically, higher Avoid accuracy is associated with smaller interference effects in the Simon task. Furthermore, as predicted, this correlation is mainly driven by increased cost of response times for incongruent trials.

To further examine the validity of our findings, a computational model of our task was developed. The goal of the computational model is to provide a two-way verification, showing (1) that our predictions indeed do follow from our hypothesis; and (2) that a model where selective attention based on the competitive basal ganglia pathways produces the expected behavioral effects.

4.1 Model Implementation

The model was implemented in the ACT-R cognitive architecture (Anderson, 2007; Anderson et al., 2004; Anderson, Fincham, Qin, & Stocco, 2008), which is perhaps the most prominent architecture currently used to model cognitive phenomena. ACT-R was chosen for two reasons. First, it is an *integrated* framework (Anderson et al., 2004) where multiple cognitive mechanisms can be simulated and connected to each other in a principled way. Second, its modules have been accurately mapped to specific brain circuits, (including the basal ganglia) thus providing a natural way to translate biological assumption into modeling constraints

(Anderson et al., 2008; Borst & Anderson, 2013; Borst, Nijboer, Taatgen, van Rijn, & Anderson, 2015; Stocco & Anderson, 2008).

An ACT-R model consists of declarative knowledge (to represent task instructions, stimulus-response mappings, and stimuli) and procedural knowledge (to represent mental operations). Declarative knowledge is represented as clusters of features, referred to as *chunks*. Chunks are processed, created, and retrieved by a set of specific modules; while they are being accessed or modified, chunks are temporarily placed in module-specific *buffers*. For example, the working memory module contains an internal, limited-capacity buffer where information can be transiently stored and updated, while a retrieval module recovers memorized chunks from long-term memory storage and places them temporarily in a retrieval buffer. Procedural knowledge, on the other hand, is represented by *production rules* (or simply *productions*). Productions are matched against the contents of the buffers and, when their specific conditions are met, result in either a motor action (for example, a manual response) or in an internal cognitive operation (for example, updating the contents of working memory).

4.1.1 Competition Between the Direct and Indirect Pathways

In ACT-R, the module responsible for the selection and execution of production rules has been mapped to the basal ganglia (Anderson et al., 2008; Stocco & Anderson, 2008). This identification rests on the empirical evidence (Anderson, 2005; Anderson et al., 2008; Stocco & Anderson, 2008) as well on the similarity of the computations performed by both systems. Crucially, the basal ganglia are believed to represent an action selection system that selects actions based on their value, and updates the value of each action through reinforcement learning (Redgrave et al., 1999; Schultz et al., 1997). Similarly, the procedural module in ACT-R selects

productions on the bases on their estimated *utility*, and updates the utility $U_p(t)$ of a production p at time t according to the following rule:

$$U_p(t) = U_p(t - 1) + \alpha[R_t - U_p(t - 1)] \quad (\text{Eq. 1})$$

Where R_t is the reward delivered at time t and α is the learning rate. It is easy to see the similarity between this equation and a canonical reinforcement learning function, with U_p taking the role of an action's Q -value, and $R_t - U_p(t - 1)$ being the prediction error.

However, ACT-R does not possess any native mechanism that captures the competitive neurodynamics between the direct and indirect pathways. Such a mechanism has been recently proposed (Stocco, submitted), and consists of splitting each production rule into a pair of competing productions corresponding to the opposing activities of the direct and indirect pathways, respectively. Like the two biological pathways they are designed to mimic, the two productions within each pair have opposite and symmetrical effects. Thus, in the case of the PSS task, if a direct pathway production selects one option, its indirect pathway version instead prevents that particular option from being selected. The utilities of these competing productions are still updated according to Eq. 1, but their learning rates are controlled by two independent parameters d_1 and d_2 , which at a biological level reflect the density of D1- and D2-type receptors in the basal ganglia.

$$U_p(t) = U_p(t - 1) + d_1 \alpha[R_t - U_p(t - 1)] \text{ for direct pathway productions}$$

$$U_p(t) = U_p(t - 1) + d_2 \alpha[R_t - U_p(t - 1)] \text{ for indirect pathway productions}$$

As shown in Stocco (submitted), this distinction is in fact needed to capture the multiple findings of the PSS task. More specifically, larger values of d_1 represent greater expression of D1 receptors and result in higher Choose accuracies *and* lower Avoid accuracies. Larger values of d_2 , conversely, represent greater expression of D2 receptors and result in lower Choose accuracies *and* higher Avoid accuracies. Most importantly, the simulations of Stocco (submitted) provide a reasonable estimate of the range of values for d_1 and d_2 that give rise to the observed variability in the sample of young, healthy participants reported in Frank et al., (2007). Specifically, it was found that individual differences could be captured by varying the values of d_1 and d_2 from 0.5 to 1.5, with the special case of $d_1 = d_2 = 1.0$ representing the average group performance.

In summary, the mechanism of competitive productions provides the necessary extension to model individual differences in Avoid and Choose accuracies. According to our hypothesis, this very same mechanism, when embedded in a plausible model of the Simon task, should produce the observed correlations between individual differences in the Simon task and individual differences in the PSS task.

Figure 6. Architecture of the ACT-R model of the Simon task. Each stimulus is represented as two features, shape and location, that compete for attention. During the selective attention phase, each feature is examined by two competing productions, a “process” production that encodes the feature in working memory, and a “don’t process” production that prevents its storage in working memory. During the response selection phase, information stored in working memory is then used to bias the retrieval of one of the two parts of the Simon task instructions that represent the appropriate stimulus-response pairs (e.g., “Press Right for Circle” and “Press Left for Square”). The retrieved stimulus-response pair is then used to execute the corresponding hand response.

The architecture of the model is illustrated in Figure 6. Like all ACT-R models, its behavior results from the interplay of cognitive modules, the content of their buffers, and production rules. The cascade of events occurring when the model responds to a stimulus, however, can be roughly divided into three subsequent stages. The first stage is the *selective attention* stage, in which the stimulus is perceived and its features are analyzed. In this stage, production rules extract relevant stimulus features from the visual buffer and place them in the working memory buffer. The second stage is the *response selection* stage, in which the correct instructions for the Simon task are retrieved from long-term memory. The third stage, not illustrated in Figure 6, is the *response check phase*, in which the model has one last chance to analyze its intended response and prevent an erroneous response.

4.1.2 Selective Attention

As in previous ACT-R models of selective attention tasks (Lovett, 2005; van Maanen, van Rijn, & Borst, 2009), the proposed model uses procedural knowledge to control selective attention to stimulus features. The competition between the direct and indirect pathways occurs precisely at this stage.

When the model sees a stimulus on the screen, its shape and location are placed in the visual perception buffer for further processing. At this point, pairs of productions compete for extracting relevant features (shape and location) and placing them in the working memory buffer. Each pair of competing production contains a “process” production that encodes the feature in working memory, and a “don’t process” production that prevents its storage in working memory. The “process” productions correspond to the activity of the direct pathway, and their learning is controlled by the d_1 parameter; while the “don’t process” productions correspond to the activity

of the indirect pathway, and their learning is controlled by the d_2 parameter. Individual differences in the d_2 parameter (which affect performance in the PSS task) are therefore related to the facility to which an individual can learn to avoid processing the location of a stimulus.

Thus, the process of selective attention depends on the competition between production rules. In ACT-R, this competition is governed by two parameters: the learning rate α (Eq. 1 and 2) and the conflict resolution noise s ; the latter determines the amount of noise that is added to each production's utility during the selection process, and thus determines the degree to which the model exploits the best available actions or explores sub-optimal ones. The third and last parameter that affects selective attention is the utility bias b , which represents the initial value for the “process” productions. This parameter is motivated by the fact that humans spontaneously process both the shape and the location of a stimulus—this is, in fact, the source of interference in the Simon task. For this reason, the two “process” productions start with utility values that are larger than those of the two “don't process” productions by the quantity b .

4.1.3 Response Selection

Because the stimuli of the Simon task do not convey any information about the response, the desired response must be either maintained in working memory or retrieved from long-term memory (Diamond, 2013). In fact, the few brain imaging studies of the Simon task have consistently reported significant regional activations in areas associated with long-term memory storage (such as the superior parietal lobule: Liu et al., 2004) or retrieval (such as the inferior frontal gyrus: Fan, Flombaum, McCandliss, Thomas, & Posner, 2003; Liu et al., 2004). Consequently, the proposed model responds by first retrieving a chunk encoding the relevant part of the Simon task instructions from long-term memory and then executing the motor command described in the rule. For example, when presented with a “Circle” stimulus, the model first

retrieves the corresponding, instructed stimulus-response mapping (“Press *Right* for *Circle*”) and then translates the rule into a corresponding motor action (“Press the *Right* index finger”; Figure 6).

In ACT-R, every chunk has an associated scalar quantity called *activation*, and the chunk with the highest activation is the one that is retrieved. The retrieval of a contextually relevant chunk is facilitated by *spreading activation* (Anderson, 1983; Collins & Loftus, 1975; Dell, 1986; Roelofs, 1992). In the model, activation spreads from the working memory buffer to the chunks in long-term memory (Anderson, 2007; Anderson et al., 2004), and is visually represented by a dashed arrow in Figure 6. The amount of activation spread is proportional to the number of features that a chunk in long-term memory shares with the contents of the buffer (Anderson & Reder, 1999; Anderson, 1990). Importantly, the amount of activation spreading from working memory is different between congruent and incongruent trials. For example, in a congruent trial both “Circle” (shape) and “Right” (location) are placed in working memory, and activation can only spread to the stimulus-response mapping “Press *Right* for *Circle*”, facilitating its retrieval (Figure 6). In incongruent trials, on the other hand, shape and location spread activation to *different* stimulus-response mappings, thus making the retrieval of the correct rule longer and more error-prone.

This difference in spreading activation is partially responsible for the difference in response times between congruent and incongruent trials. This is because, in ACT-R, the retrieval time RT_c for a chunk c is related to a chunk’s activation $A(c)$ by the following equation:

$$RT_c = F e^{-fA(c)} \quad (\text{Eq. 3})$$

Because the two parameters in Eq. 3, F and f , have similar effects on the model's behavior, many models of declarative memory fit only one of them, leaving the other one to its default value. This is the case, for example, of John Anderson's influential model of the fan effect (Anderson & Reder, 1999), which is currently included as a reference model in the ACT-R distribution. In our case, we chose to examine the effects of F on the model's behavior, and leave f to its default value of 1.0. A second parameter n , determines the amount of noise in the activation system, which plays a role akin to s for the production's utility.

4.1.4 Response Monitoring

After a response has been selected, the model performs an internal check (not illustrated in Figure 6) of the consistency of the rule with the stimulus' shape. If the retrieved rule does not match the current stimulus shape (for example, the retrieved rule is "if *Circle*, press *Right*", but the stimulus is a square), the model scratches the contents of working memory and re-processes the stimulus; otherwise, it proceeds to respond. This internal monitoring has the purpose of delivering an internal feedback signal when an incorrect response is detected, which is in turn needed for learning. Because incongruent trials are more likely to result in the retrieval of the wrong stimulus-response mappings, they are also more likely to incur in a restart of the retrieval process, further compounding the incongruency effect.

4.1.5 Sensorimotor Processes

ACT-R also provides a realistic simulation of the sensorimotor processes of the task. In the model, sensorimotor processes are only needed for encoding a stimulus and performing a motor response. The time T to encode a new stimulus on the screen is set to 185 ms, which is the unmodified default value in ACT-R. The motor response time M in ACT-R is the sum of the times needed to process and execute a movement, which amounts to 100ms. The model assumes

that, given the simplicity and the repetitiveness of the two motor responses, no additional preparation time is needed. Note that both T and M are fixed, to their default values and, since they cannot affect the model's predicted patterns (i.e., the relative response times and accuracies for congruent and incongruent trials), and can be considered as a single constant quantity.

4.1.6 Model Parameters

Table 1 provides a list of the parameters used in the model, with an indication of the range under which their effect was examined, their final value, and their function. The parameter's *value* refers to the set of parameters that yield the best fit and whose results are illustrated below. The parameter's *range* refers to the range of values over which a full grid search was conducted to explore the model behavior and evaluate the model performance.

Table 1: Parameters used in the ACT-R model of the Simon task (Figure 6).

Name	Function	Range	Value
<i>Procedural memory (Selective attention)</i>			
α	Utility learning rate (Eq 1 and 2)	0. - 0.5	0.25
s	Utility noise	0.0 - 1.0	0.0
b	Bias for “process” productions	1.0 - 10.0	10.0
<i>Declarative memory (Response selection)</i>			
F	Retrieval latency factor (Eq. 3)	0.2 - 1.0	0.25
n	Activation noise	0.0 - 1.0	0.2
f	Retrieval exponent factor (Eq. 3)	Fixed	1.0
<i>Sensorimotor processes</i>			
T	Stimulus encoding time	Fixed	185 ms
M	Motor execution time	Fixed	100 ms
<i>Correlations between Simon and PSS</i>			
d_1	D1 receptor density (Eq. 2)	0.5-1.5	1
d_2	D2 receptor density (Eq. 2)	0.5-.15	1

4.3. Simulations

To properly evaluate a model, it is important to examine its *fit* to the experimental data under optimal conditions as well as its *complexity*, that is, the degrees of freedom allowed by its parameters (Roberts & Pashler, 2000). Both fit and complexity can be evaluated by performing a full grid search over the parameter space. Because such a grid search is a computationally expensive procedure, the range of values over which a parameter is evaluated needs to be principled. This can be done by constraining the parameter range within the limits under which it

produces meaningful behavior. Although all of the parameters in Figure 1 can *in principle* take any real value, negative values are meaningless for all of them. For instance, a negative value of α would result in unlearning. Similarly, above a certain positive value, the effect of a parameter on a model effectively stops having any meaningful effect. For example, for large values of s , the model effectively selects at random between competing productions; and, for large values of n , it selects at random between competing chunks. The upper values in Table 1 were chosen on the basis of these considerations.

4.4 Model Fit to Experimental Data

In performing the grid search procedure, the values of α , s , F , n , and b were fully crossed, but the different values of d_1 and d_2 were examined separately. Thus, for each combination of values of α , s , F , n , and b , the model was then run 100 times for each value of d_1 in the range of d_1 while maintaining $d_2 = 1.0$, and then 100 times for each value in the range of d_2 while maintaining $d_1 = 1.0$.

The best fit to the group results for the Simon task was found by identifying the hyperpoint in the parameter space that minimized the sum of squared errors between the predicted and the actual mean responses times and accuracies. Because this analysis concerns a group average, the values of d_1 and d_2 were set to 1.0. Under these conditions, both equations in Eq. 2 reduce to the standard form of Eq. 1, and the “process” and “don’t process” productions are in equilibrium. The optimal values for each parameter are reported in Table 1. Figure 7 visually illustrates the model predictions against the group results for the Simon task. The model captures both the main effects of condition on accuracy ($RMSE = 0.02$, $\chi^2(1) = 0.0003$, $p > 0.98$; Figure 7A) and response times ($RMSE = 19.02$ ms, $\chi^2(1) = 0.00002$, $p = 0.996$; Figure 7B).

Figure 7: Experimental results (light grey) and model predictions (dark grey) for the Simon Task under the best fitting parameter values (Table 1). (A) Accuracy results (B) Response times. Bars represent mean values. Error bars represent standard errors of the mean.

For the purpose of the present study, however, it is crucial to also examine how varying values of d_1 and d_2 affect the response times in the Simon task. These results are represented in Figure 8. Because the pattern of correlations is intrinsically more variable than the average results in the Simon task (see the next section for a full analysis of this variability), the results presented in Figure 8 were obtained by maintaining the F and n parameters at their optimal value (see Table 1) and averaging the model's performance across all the levels of the α , s and b . This was done because the F and n parameters significantly affect the response times (by orders of magnitude), clouding the comparison with Figure 7. The results in Figure 8 are thus

representative of what can be expected in terms of individual differences while maintaining the mean response times in lines with the results of Figure 7.

As shown in Figure 8B, and consistent with our hypothesis, response times for incongruent trials decline sharply with increasing values of d_2 , effectively reducing the Simon effect by more than half (from 40 ms to 17ms). On the other hand, the response times for congruent trials remained essentially identical, with a difference of only ~2 ms between the minimum and the maximum values. The effects of the d_1 parameter on Simon task performance reveal a substantially different pattern (Figure 8A). As in the case of d_2 , response times for congruent trials did not correlate with the values of d_1 . In the case of incongruent trials, however, increased d_1 resulted in slightly *longer* (8 ms) incongruent response times for larger values of d_1 .

In summary, the model predicts that increased amount of D2 receptors would result in smaller incongruency effects in the Simon task and smaller response times during incongruent trials (Figure 8). Furthermore, the model also predicts that response times for congruent trials should not be significantly affected by the individual differences in D1 receptors, with possibly only a small positive increase in incongruent response times for larger values of D1.

Figure 8: Correlations between the values of parameter d_1 (A) and parameter d_2 (B) and the predicted response times for the Simon task under the best-fitting values of F and n . Grey lines and points represent the response times for congruent and incongruent trials, averaged over 100 runs of the model; dotted lines represent best-fitting regression lines.

4.5 Model Complexity

Although the best fit of a model provides an intuitive measure of its quality, model evaluation should also consider the model's *complexity* (Roberts & Pashler, 2000), with the assumption that models should be penalized for excessive complexity. Note that the raw number of parameters in a model provides a poor estimation of complexity, as multiple parameters do not necessarily entail a greater probability of fitting the data (Veksler, Myers, & Gluck, 2015). Model complexity can instead be inferred by examining how model performance depends on the full set of parameters (Pitt, Kim, Navarro, & Myung, 2006; Veksler et al., 2015). Intuitively, a

model that produces all the possible outcomes of an experiment should be penalized more than a model that predicts only a few selected outcomes.

This complexity can be characterized through a method known as *parameter space partitioning* (PSP), in which the performance of a model is analyzed at each point in model parameter space, and the size of regions corresponding to different data patterns is then measured. Intuitively, a model that reproduces the pattern of results of the experimental data in most cases is superior to a model that can only produce the same pattern under very specific sets of parameters. In examining our model, two types of data space need to be kept in mind. One is the data pattern corresponding to the raw Simon task data. The other is the data pattern corresponding to the correlations between the Simon effects and the PSS performance.

Data from the Simon task can yield nine possible data patterns, depending on whether accuracies and response times are greater for congruent trials, incongruent trials, or equal (“equal” was counted as any pattern whereby the difference between the smaller and larger values was less than 2%). An analysis of the model’s performance across the full parameter space revealed that, out of these nine theoretically possible patterns, only three were produced by the model. In all of these patterns, accuracies were higher for congruent than for incongruent trials. Thus, the model never predicts equal accuracies for both types of trials, nor does it predict the opposite pattern of greater accuracy for incongruent trials. Of the three patterns predicted by the model, the correct pattern (with greater response times for incongruent than for congruent trials) accounted for 99.8% of all possible results. The remaining two patterns, each of which accounts for 0.1% of all the possible cases, correspond to the cases in which the model either predicts no difference in response times between the two conditions or predicts greater response times for the congruent trials. These results are visually summarized in Figure 9A.

Parameter space partitioning can be similarly used to assess the flexibility of the model in predicting the pattern of correlations between the Simon task and the d_1 and d_2 values. In this case, nine patterns are theoretically possible based on the sign (positive, negative, or zero) of the correlations between the value of the d_2 parameter (which predicts Avoid accuracy) and the value of the response times for incongruent trials, as well as the correlation between d_2 and the Simon effect. In this case, all nine possible patterns occur as model predictions, but with very different frequencies. Specifically, the correct result pattern, whereby the model predicts lower r values for d_2 correlations, occurs 76.0% of the time, while the remaining eight possible pattern make up the remaining 24.0% of the cases.

Figure 9: Results of Parameter Space Partitioning (PSP) analysis. (A) PSP analysis of the group averages in the Simon task; the model consistently predicts the correct pattern, with incongruent trials yielding lower accuracies and higher response times than congruent trials; (B) PSP analysis of correlations between two indices of Simon task performance and d_2 parameter. The model predominantly predicts negative correlations of d_2 with both the Simon effect and the mean response for incongruent trials.

5. Conclusions

This paper has presented the hypothesis that one feature of cognitive control, that is, the allocation of attention to specific dimensions of the stimuli, is related to the activity of the basal ganglia, and especially to the circuit's indirect pathway. This hypothesis was both tested experimentally and examined in detail through a computational model. Experimentally, a correlation was found between Avoid accuracy in the PSS task and the size of the interference

effect in the Simon task. As expected, this correlation was driven by the additional cost of incongruent trials, rather than by savings during congruent trials, and was specific to the indirect pathway (as evidenced by a lack of a similar negative correlation between Choose accuracy and response times in the Simon task). Furthermore, this correlation was found to be independent of individual differences in working memory capacity. Additionally, computer simulations showed that an ACT-R model of the Simon task, properly modified to include the competitive dynamics of the two basal ganglia pathways, does indeed predict the same pattern of results that is observed in our data. Furthermore, a parameter space partitioning examination of the sensitivity of the model's behaviors to its parameters value showed that the model's predictions are remarkably robust and constrained across very large sections of the parameter space.

5.1 Relationship to Other Models

The proposed model predicts that individual differences in cognitive control are at least partially related to individual differences in the differential learning rates of the two pathways in the basal ganglia. In doing so, the model draws a connection between two apparently separate modeling traditions, models of cognitive control and models of reinforcement learning.

Although novel, the proposed model is not in contrast with previous models of cognitive control. In fact, like many previous models of cognitive control (Cohen et al., 1990), our model also suppresses prepotent but inappropriate responses through the use of spreading activation from internal task representations. In this specific case, activation spreads from the working memory buffer (Figure 6), whose function is closely related to that of PFC. Differently than in previous models, however, this internal representation is not directly provided by the author in the form of a fixed input but is internally generated and revised by the model itself. Specifically, the model relies on an action selection system to learn which features should be included in the

representation. This action selection system represents the functional equivalent of the basal ganglia (Anderson et al., 2008; Stocco & Anderson, 2008; Stocco et al., 2010), and uses reinforcement learning to refine its choices on the bases of its monitored behavioral performance. Thus, our proposed model does not do away with the consensus view of cognitive control as a top-down mechanism based on internal task representations; rather, it supplements this view with a bottom-up mechanism that adaptively chooses and refines those internal representations. According to this view, cognitive control arises in the dynamic interplay between the basal ganglia and prefrontal cortex, rather than as a static function of prefrontal cortex alone.

The model learns how to shape its internal task representations through reinforcement learning (Sutton & Barto, 1998), a framework for unsupervised learning algorithms that has been often used to interpret the function of the basal ganglia (Frank et al., 2004; Joel, Niv, & Ruppín, 2002/6; O'Reilly & Frank, 2006; Schultz et al., 1997; Stocco, 2012). Despite their simple premises, reinforcement learning models have been recently demonstrated to be capable of producing remarkably complex behaviors (Mnih et al., 2015; Silver et al., 2016). This is particularly true of models of *hierarchical* reinforcement learning (Botvinick, 2012; Botvinick, Niv, & Barto, 2009; Collins & Frank, 2013), which can perform learn complex tasks by creating “options”, that is, structured internal representations that control action executions for long periods of time. Several authors have suggested that, while the basal ganglia might execute actions and learn their value, PFC seems ideally positioned to maintain and represent options, which in turn control the execution of future actions (Badre & Frank, 2012; Botvinick, 2012; Frank & Badre, 2012). In a sense, our model can be seen as a very simple form of hierarchical reinforcement learning, whereby the internal representations in the working memory buffer

represent a minimalistic option, and the model learns to select actions that ultimately shape the available option.

A unique facet of our model is that, in addition to reproducing the average group results on the Simon task, it also provides a mechanistic explanation for individual differences in Simon task performance. The same mechanism and parameter values were also used to explain individual differences on the PSS task itself in a different model (Stocco, submitted). Although models exist that perform the Simon task (Dutt, Yamaguchi, Gonzalez, & Proctor, 2009) or similar interference tasks (Cohen et al., 1990; Lovett, 2005; van Maanen et al., 2009), and the PSS task (Frank et al., 2004) separately, ours is the only case of a model designed to account for both.

The integration of top-down control and reinforcement learning, together with the use of the same code and parameters across two different tasks, is made possible by the use of a single cognitive architecture—specifically, ACT-R. ACT-R has been criticized in the past for allowing too many degrees of flexibility (Roberts & Pashler, 2000); while this criticism is legitimate in the many cases of single, one-off models, ACT-R was explicitly designed for the purpose of facilitating model re-usability and integration. We believe that the model described herein represents a good case for the validity of this approach.

5.2 Limitations and Future Directions

Both our results and our model crucially depend on a number of assumptions. Perhaps the most important is that Choose and Avoid accuracies in the PSS task do adequately reflect the activity of the direct and indirect pathway. Although previous studies have confirmed that experimental manipulations of dopamine levels (Frank et al., 2004) and genetic expression of D1 and D2 receptors (Frank et al., 2007) affect these two measures, the degree to which Choose and

Avoid accuracies are predictive of the underlying biology remain to be empirically verified. It is possible, in principle, that both measures are partially confounded by other variables, such as individual differences in working memory or the choice of specific strategies. However, the PSS was explicitly designed to discourage participants from relying on “explicit” strategies that would depend on working memory; for example, it is probabilistic in nature and employs stimuli that are unfamiliar and difficult to verbalize (Frank et al., 2004). In this sense, it precisely mimics the characteristics of tasks that have been successfully used to study procedural learning in the basal ganglia in humans (Knowlton & Squire, 1994). Additionally, although participants might still use idiosyncratic strategies, the existence of a performance criterion that must be passed before concluding the training phase provides a safeguard against the most trivial ones (e.g., “always select the stimulus on the right”) that might confound the data. Finally, it should be noted that no better measure currently exists to non-invasively characterize the relative strengths of the pathways; even genetic studies are limited to considering only a narrow group of known polymorphisms (Frank et al., 2007; Zhang et al., 2007), and, due to the characteristics of the BOLD signal, the relative contributions of the two pathways cannot be reconstructed from fMRI data (Aron, 2007; Hershey et al., 2003).

Another underlying assumption is that the interference effect in the Simon task is driven by selective attention failure; more specifically, by a failure to avoid processing the irrelevant stimulus dimension of position. This assumption might seem incompatible with data suggesting that interference in the Simon task occurs at the level of the response, when both hands are activated, rather than at the level of stimulus encoding (Liu et al., 2004; Nee et al., 2007). The incompatibility is only apparent because the interference at the response level can only occur if the irrelevant dimension (position) has been processed earlier—otherwise, there is no reason to

believe that the irrelevant hand response could be activated. Indeed, interference at the response level does occur in our model as well, at the moment when competing S-R mappings are retrieved. In other words, interference might manifest itself at a later stage of the task, but have its remote cause in processing occurring earlier on. The model by van Maanen and colleagues (2009), for example, provides an elegant account of how interference in two apparently different tasks (picture-word naming and Stroop) might originate at the same early stage from competing stimulus-encoding processes, but manifest at different later stages because of the nature and characteristic durations of the specific processes involved (van Maanen et al., 2009).

Another limitation of the current study is that cognitive control data was collected from only one behavioral task. The Simon task was chosen because it is behaviorally simple, reliable, commonly adopted in cognitive control research, and it involves only memory and perceptual processes that are rather uncontroversial. In our opinion, these characteristics make the Simon task a better choice than, for instance, the Stroop task, whose execution also crucially depends on linguistic processes. However, both the Simon task and the Stroop task pose similar demands on selective attention, and could potentially be similarly related to individual differences in the PSS task.

In the current study, a decision was made to include a 3:1 ratio of incongruent to congruent trials. Asymmetric ratios have been commonly used in studying individual differences in cognitive control (e.g., Long & Prat, 2002), likely because they amplify the difference between the two conditions and provide greater variability across individuals. While the smaller frequency of the incongruent condition might account for part of our effects, it is unlikely to be the main source. For example, the so-called oddball paradigms, which rely solely on the different frequencies of two classes of stimuli, typically employ much more extreme ratios (often on the

order of 10:1, e.g. García-Larrea, Lukaszewicz, & Mauguière, 1992) to induce observable responses—even in the form of event-related potentials, which are arguably more sensitive measures than response times.

A final limitation concerns the set of possible alternative biological mechanisms that might underlie our experimental findings. In this paper, we have examined the evidence for the role of the basal ganglia's direct and indirect pathways in a cognitive control task, and we were able to rule out any the possible confounds of working memory capacity. However, other possible biological mechanisms exist that could conceivably explain our data. For example, the basal ganglia also contain a hyperdirect pathway that, much like the indirect pathway, plays an important role in inhibiting unwanted signals (Frank, Samanta, et al., 2007; Nambu, Tokuno, & Takada, 2002).

5.3 Summary

These limitations notwithstanding, we believe that our results provide new insights into the biological underpinnings of cognitive control, and suggest new ways to investigate this topic. Specifically, our findings suggest a framework whereby cognitive control originates from the interplay of basal ganglia and prefrontal cortex. Furthermore, they also suggest that different components of cognitive control might correspond to the different directions of this dynamic exchange. Specifically, our results imply that selective attention is related to the role of the basal ganglia in selecting the information to be encoded in PFC. We see this as a starting point for future investigations of the nature of cognitive control and individual differences in cognitive abilities.

6. Supplementary Material

All of the experimental and modeling datasets (including the individual subject log files, the aggregated tables, the ACT-R code for the model Figure 6, additional Lisp and Python code to run the simulations, and all the data from the simulations) and the R code to analyze the data are freely available at our laboratory's repository: https://github.com/UWCCDL/PSS_Simon/.

7. Acknowledgments

This research was funded by a grant from the Office of Naval Research (ONRBAA13-003) entitled “Training the Mind and Brain: Investigating Individual Differences in the Ability to Learn and Benefit Cognitively from Language Training.”

8. References

- Albin, R. L., Young, A. B., & Penney, J. B. (1989). The functional anatomy of basal ganglia disorders. *Trends in Neurosciences*, 12(10), 366–375.
- Alexander, G. E., DeLong, M. R., & Strick, P. L. (1986). Parallel organization of functionally segregated circuits linking basal ganglia and cortex. *Annual Review of Neuroscience*, 9, 357–381.
- Anderson, J. R. (1983). A spreading activation theory of memory. *Journal of Verbal Learning and Verbal Behavior*, 22(3), 261–295.
- Anderson, J. R. (1990). *The Adaptive Character of Thought*. Mahwah, NJ: Lawrence Erlbaum Associates.
- Anderson, J. R. (2005). Human symbol manipulation within an integrated cognitive architecture. *Cognitive Science*, 29(3), 313–341.

- Anderson, J. R. (2007). *How Can the Mind Occur in the Physical Universe?* New York, NY: Oxford University Press.
- 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.
- Anderson, J. R., Fincham, J. M., Qin, Y., & Stocco, A. (2008). A central circuit of the mind. *Trends in Cognitive Sciences*, *12*(4), 136–143.
- Anderson, J. R., & Reder, L. M. (1999). The fan effect: New results and new theories. *Journal of Experimental Psychology: General*, *128*(2), 186.
- Aron, A. R. (2007). The neural basis of inhibition in cognitive control. *The Neuroscientist: A Review Journal Bringing Neurobiology, Neurology and Psychiatry*, *13*(3), 214–228.
- Awh, E., Vogel, E. K., & Oh, S.-H. (2006). Interactions between attention and working memory. *Neuroscience*, *139*(1), 201–208.
- Badre, D., & Frank, M. J. (2012). Mechanisms of hierarchical reinforcement learning in cortico-striatal circuits 2: Evidence from fMRI. *Cerebral Cortex*, *22*(3), 527–536.
- Berman, S. M., & Noble, E. P. (1995). Reduced visuospatial performance in children with the D2 dopamine receptor A1 allele. *Behavior Genetics*, *25*(1), 45–58.
- Borst, J. P., & Anderson, J. R. (2013). Using model-based functional MRI to locate working memory updates and declarative memory retrievals in the fronto-parietal network. *Proceedings of the National Academy of Sciences of the United States of America*, *110*(5), 1628–1633.

- Borst, J. P., Nijboer, M., Taatgen, N. A., van Rijn, H., & Anderson, J. R. (2015). Using data-driven model-brain mappings to constrain formal models of cognition. *PLOS One*, *10*(3), e0119673.
- Botvinick, M. M. (2012). Hierarchical reinforcement learning and decision making. *Current Opinion in Neurobiology*, *22*(6), 956–962.
- Botvinick, M. M., Niv, Y., & Barto, A. C. (2009). Hierarchically organized behavior and its neural foundations: a reinforcement learning perspective. *Cognition*, *113*(3), 262–280.
- Botvinick, M., & Plaut, D. C. (2004). Doing without schema hierarchies: a recurrent connectionist approach to normal and impaired routine sequential action. *Psychological Review*, *111*(2), 395–429.
- Braver, T. S. (2012). The variable nature of cognitive control: A dual mechanisms framework. *Trends in Cognitive Sciences*, *16*(2), 106–113.
- Braver, T. S., Paxton, J. L., Locke, H. S., & Barch, D. M. (2009). Flexible neural mechanisms of cognitive control within human prefrontal cortex. *Proceedings of the National Academy of Sciences of the United States of America*, *106*(18), 7351–7356.
- Cohen, J. D., Braver, T. S., & O'Reilly, R. C. (1996). A computational approach to prefrontal cortex, cognitive control and schizophrenia: Recent developments and current challenges. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, *351*(1346), 1515–1527.
- Cohen, J. D., Dunbar, K., & McClelland, J. L. (1990). On the control of automatic processes: a parallel distributed processing account of the Stroop effect. *Psychological Review*, *97*(3),

332–361.

- Cole, M. W., Laurent, P., & Stocco, A. (2013). Rapid instructed task learning: A new window into the human brain's unique capacity for flexible cognitive control. *Cognitive, Affective & Behavioral Neuroscience*, 13(1), 1–22.
- Cole, M. W., & Schneider, W. (2007). The cognitive control network: Integrated cortical regions with dissociable functions. *NeuroImage*, 37(1), 343–360.
- Collins, A. G. E., & Frank, M. J. (2013). Cognitive control over learning: creating, clustering, and generalizing task-set structure. *Psychological Review*, 120(1), 190–229.
- Collins, A. M., & Loftus, E. F. (1975). A spreading-activation theory of semantic processing. *Psychological Review*, 82(6), 407.
- Craft, J. L., & Simon, J. R. (1970). Processing symbolic information from a visual display: interference from an irrelevant directional cue. *Journal of Experimental Psychology*, 83(3), 415–420.
- Dell, G. S. (1986). A spreading-activation theory of retrieval in sentence production. *Psychological Review*, 93(3), 283–321.
- DeLong, M. R. (1990). Primate models of movement disorders of basal ganglia origin. *Trends in Neurosciences*, 13(7), 281–285.
- Diamond, A. (2013). Executive functions. *Annual Review of Psychology*, 64, 135–168.
- Dutt, V., Yamaguchi, M., Gonzalez, C., & Proctor, R. W. (2009). An instance-based learning model of stimulus-response compatibility effects in mixed location-relevant and location-

irrelevant tasks. In *Proceedings of the 9th International Conference on Cognitive Modeling—ICCM2009*, Manchester, UK.

Engle, R. W., Kane, M. J., & Tuholski, S. W. (1999). Individual differences in working memory capacity and what they tell us about controlled attention, general fluid intelligence, and functions of the prefrontal cortex. In A. Miyake & P. Shah (Eds). *Models of Working memory: Mechanisms of Active maintenance and Executive Control*, (pp. 102-134). New York, NY, US: Cambridge University Press

Fan, J., Flombaum, J. I., McCandliss, B. D., Thomas, K. M., & Posner, M. I. (2003). Cognitive and brain consequences of conflict. *NeuroImage*, 18(1), 42–57.

Frank, M. J., & Badre, D. (2012). Mechanisms of hierarchical reinforcement learning in corticostriatal circuits 1: computational analysis. *Cerebral Cortex*, 22(3), 509–526.

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.

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 of the United States of America*, 104(41), 16311–16316.

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

- Frank, M. J., Seeberger, L. C., & O'reilly, R. C. (2004). By carrot or by stick: Cognitive reinforcement learning in parkinsonism. *Science*, 306(5703), 1940–1943.
- García-Larrea, L., Lukaszewicz, A. C., & Mauguière, F. (1992). Revisiting the oddball paradigm. Non-target vs neutral stimuli and the evaluation of ERP attentional effects. *Neuropsychologia*, 30(8), 723–741.
- Gerfen, C. R., Engber, T. M., Mahan, L. C., Susel, Z., Chase, T. N., Monsma, F. J., Jr, & Sibley, D. R. (1990). D1 and D2 dopamine receptor-regulated gene expression of striatonigral and striatopallidal neurons. *Science*, 250(4986), 1429–1432.
- Gulbinaite, R., van Rijn, H., & Cohen, M. X. (2014). Fronto-parietal network oscillations reveal relationship between working memory capacity and cognitive control. *Frontiers in Human Neuroscience*, 8, 761.
- Gurney, K., Prescott, T. J., & Redgrave, P. (2001). A computational model of action selection in the basal ganglia. I. A new functional anatomy. *Biological Cybernetics*, 84(6), 401–410.
- Han, D. H., Yoon, S. J., Sung, Y. H., Lee, Y. S., Kee, B. S., Lyoo, I. K., ... Cho, S. C. (2008). A preliminary study: novelty seeking, frontal executive function, and dopamine receptor (D2) TaqI A gene polymorphism in patients with methamphetamine dependence. *Comprehensive Psychiatry*, 49(4), 387–392.
- Herd, S. A., Banich, M. T., & O'Reilly, R. C. (2006). Neural mechanisms of cognitive control: an integrative model of stroop task performance and fMRI data. *Journal of Cognitive Neuroscience*, 18(1), 22–32.
- Hershey, T., Revilla, F. J., Wernle, A. R., McGee-Minnich, L., Antenor, J. V., Videen, T. O., ...

- Perlmutter, J. S. (2003). Cortical and subcortical blood flow effects of subthalamic nucleus stimulation in PD. *Neurology*, *61*(6), 816–821.
- Houk, J. C., Bastianen, C., Fansler, D., Fishbach, A., Fraser, D., Reber, P. J., ... Simo, L. S. (2007). Action selection and refinement in subcortical loops through basal ganglia and cerebellum. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, *362*(1485), 1573–1583.
- Humphries, M. D., Stewart, R. D., & Gurney, K. N. (2006). A physiologically plausible model of action selection and oscillatory activity in the basal ganglia. *The Journal of Neuroscience*, *26*(50), 12921–12942.
- Joel, D., Niv, Y., & Ruppín, E. (2002/6). Actor–critic models of the basal ganglia: new anatomical and computational perspectives. *Neural Networks*, *15*(4–6), 535–547.
- Kane, M. J., & Engle, R. W. (2002). The role of prefrontal cortex in working-memory capacity, executive attention, and general fluid intelligence: An individual-differences perspective. *Psychonomic Bulletin & Review*, *9*(4), 637–671.
- Kane, M. J., & Engle, R. W. (2003). Working-memory capacity and the control of attention: the contributions of goal neglect, response competition, and task set to Stroop interference. *Journal of Experimental Psychology: General*, *132*(1), 47–70.
- Kemp, J. M., & Powell, T. P. (1971). The site of termination of afferent fibres in the caudate nucleus. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, *262*(845), 413–427.
- Koechlin, E., Ody, C., & Kouneiher, F. (2003). The architecture of cognitive control in the

- human prefrontal cortex. *Science*, 302(5648), 1181–1185.
- Liu, X., Banich, M. T., Jacobson, B. L., & Tanabe, J. L. (2004). Common and distinct neural substrates of attentional control in an integrated Simon and spatial Stroop task as assessed by event-related fMRI. *NeuroImage*, 22(3), 1097–1106.
- Long, D. L., & Prat, C. S. (2002). Working memory and Stroop interference: an individual differences investigation. *Memory & Cognition*, 30(2), 294–301.
- Lovett, M. C. (2005). A Strategy-based interpretation of Stroop. *Cognitive Science*, 29(3), 493–524.
- Lu, C. H., & Proctor, R. W. (1995). The influence of irrelevant location information on performance: A review of the Simon and spatial Stroop effects. *Psychonomic Bulletin & Review*, 2(2), 174–207.
- MacDonald, A. W., Cohen, J. D., Stenger, V. A., & Carter, C. S. (2000). Dissociating the role of the dorsolateral prefrontal and anterior cingulate cortex in cognitive control. *Science*, 288(5472), 1835–1838.
- McNab, F., & Klingberg, T. (2008). Prefrontal cortex and basal ganglia control access to working memory. *Nature Neuroscience*, 11(1), 103–107.
- Messer, C. (2011). Examining the influence of the dopamine D2 agonist, cabergoline, on measures of response inhibition and error awareness. *Retrieved from* <http://espace.library.uq.edu.au/view/UQ:275915>
- Miller, E. K. (2000). The prefrontal cortex and cognitive control. *Nature Reviews Neuroscience*,

I(1), 59–65.

Miller, E. K., & Cohen, J. D. (2001). An integrative theory of prefrontal cortex function. *Annual Review of Neuroscience*, 24, 167–202.

Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., ... Hassabis, D. (2015). Human-level control through deep reinforcement learning. *Nature*, 518(7540), 529–533.

Nambu, A., Tokuno, H., & Takada, M. (2002). Functional significance of the cortico-subthalamo-pallidal “hyperdirect” pathway. *Neuroscience Research*, 43(2), 111–117.

Nee, D. E., Wager, T. D., & Jonides, J. (2007). Interference resolution: insights from a meta-analysis of neuroimaging tasks. *Cognitive, Affective & Behavioral Neuroscience*, 7(1), 1–17.

Niv, Y., Duff, M. O., & Dayan, P. (2005). Dopamine, uncertainty and TD learning. *Behavioral and Brain Functions*, 1, 6.

O'Reilly, R. C., & Frank, M. J. (2006). Making working memory work: a computational model of learning in the prefrontal cortex and basal ganglia. *Neural Computation*, 18(2), 283–328.

Pitt, M. A., Kim, W., Navarro, D. J., & Myung, J. I. (2006). Global model analysis by parameter space partitioning. *Psychological Review*, 113(1), 57–83.

Proctor, R. W., & Lu, C. H. (1999). Processing irrelevant location information: practice and transfer effects in choice-reaction tasks. *Memory & Cognition*, 27(1), 63–77.

Rainer, G., Asaad, W. F., & Miller, E. K. (1998). Selective representation of relevant information by neurons in the primate prefrontal cortex. *Nature*, 393(6685), 577–579.

- Redgrave, P., Prescott, T. J., & Gurney, K. (1999). The basal ganglia: A vertebrate solution to the selection problem? *Neuroscience*, 89(4), 1009–1023.
- Ridderinkhof, K. R., van den Wildenberg, W. P. M., Segalowitz, S. J., & Carter, C. S. (2004). Neurocognitive mechanisms of cognitive control: The role of prefrontal cortex in action selection, response inhibition, performance monitoring, and reward-based learning. *Brain and Cognition*, 56(2), 129–140.
- Roberts, S., & Pashler, H. (2000). How persuasive is a good fit? A comment on theory testing. *Psychological Review*, 107(2), 358–367.
- Roelofs, A. (1992). A spreading-activation theory of lemma retrieval in speaking. *Cognition*, 42(1-3), 107–142.
- Schultz, W. (1998). Predictive reward signal of dopamine neurons. *Journal of Neurophysiology*, 80(1), 1–27.
- Schultz, W., Apicella, P., & Ljungberg, T. (1993). Responses of monkey dopamine neurons to reward and conditioned stimuli during successive steps of learning a delayed response task. *The Journal of Neuroscience*, 13(3), 900–913.
- Schultz, W., Dayan, P., & Montague, P. R. (1997). A neural substrate of prediction and reward. *Science*, 275(5306), 1593–1599.
- Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., van den Driessche, G., ... Hassabis, D. (2016). Mastering the game of Go with deep neural networks and tree search. *Nature*, 529(7587), 484–489.

- Simon, J. R. (1990). The Effects of an Irrelevant Directional CUE on Human Information Processing. In Robert W. Proctor and T. Gilmour Reeve (Ed.), *Advances in Psychology* (Vol. Volume 65, pp. 31–86). North-Holland.
- Stephenson-Jones, M., Samuelsson, E., Ericsson, J., Robertson, B., & Grillner, S. (2011). Evolutionary conservation of the basal ganglia as a common vertebrate mechanism for action selection. *Current Biology*, 21(13), 1081–1091.
- Stewart, T. C., Choo, X., & Eliasmith, C. (2010). Dynamic behaviour of a spiking model of action selection in the basal ganglia. In *Proceedings of the 10th international conference on cognitive modeling* (pp. 235–240).
- Stocco, A. (2012). Acetylcholine-based entropy in response selection: a model of how striatal interneurons modulate exploration, exploitation, and response variability in decision-making. *Frontiers in Neuroscience*, 6, 18.
- 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.
- Stocco, A., & Lebiere, C. (2014). Inhibitory synapses between striatal projection neurons support efficient enhancement of cortical signals: A computational model. *Journal of Computational Neuroscience*, 37(1), 65–80.
- 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.
- Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning: An introduction*. Cambridge, MA:

MIT Press.

Unsworth, N., & Engle, R. W. (2007). On the division of short-term and working memory: an examination of simple and complex span and their relation to higher order abilities.

Psychological Bulletin, 133(6), 1038–1066.

Unsworth, N., Heitz, R. P., Schrock, J. C., & Engle, R. W. (2005). An automated version of the operation span task. *Behavior Research Methods*, 37(3), 498–505.

van Maanen, L., van Rijn, H., & Borst, J. P. (2009). Stroop and picture—word interference are two sides of the same coin. *Psychonomic Bulletin & Review*, 16(6), 987–999.

van Veen, V., & Carter, C. S. (2005). Separating semantic conflict and response conflict in the Stroop task: a functional MRI study. *NeuroImage*, 27(3), 497–504.

Veksler, V. D., Myers, C. W., & Gluck, K. A. (2015). Model flexibility analysis. *Psychological Review*, 122(4), 755–769.

Vogel, E. K., McCollough, A. W., & Machizawa, M. G. (2005). Neural measures reveal individual differences in controlling access to working memory. *Nature*, 438(7067), 500–503.

Zhang, Y., Bertolino, A., Fazio, L., Blasi, G., Rampino, A., Romano, R., ... Sadée, W. (2007).

Polymorphisms in human dopamine D2 receptor gene affect gene expression, splicing, and neuronal activity during working memory. *Proceedings of the National Academy of Sciences of the United States of America*, 104(51), 20552–20557.