

Published in final edited form as:

Neural Netw. 2012 August ; 32: 74–85. doi:10.1016/j.neunet.2012.02.033.

Simulating the Effects of Dopamine Imbalance on Cognition: From Positive Affect to Parkinson's Disease

Sébastien Hélie, Erick J. Paul, and F. Gregory Ashby

University of California, Santa Barbara

Abstract

Cools (2006) suggested that prefrontal dopamine levels are related to cognitive stability whereas striatal dopamine levels are related to cognitive plasticity. With such a wide ranging role, almost all cognitive activities should be affected by dopamine levels in the brain. Not surprisingly, factors influencing brain dopamine levels have been shown to improve/worsen performance in many behavioral experiments. On the one hand, Nadler and her colleagues (2010) showed that positive affect (which is thought to increase cortical dopamine levels) improves a type of categorization that depends on explicit reasoning (rule-based) but not a type that depends on procedural learning (information-integration). On the other hand, Parkinson's disease (which is known to decrease dopamine levels in both the striatum and cortex) produces proactive interference in the odd-man-out task (Flowers & Robertson, 1985) and renders subjects insensitive to negative feedback during reversal learning (Cools et al., 2006). This article uses the COVIS model of categorization to simulate the effects of different dopamine levels in categorization, reversal learning, and the odd-man-out task. The results show a good match between the simulated and human data, which suggests that the role of dopamine in COVIS can account for several cognitive enhancements and deficits related to dopamine levels in healthy and patient populations.

Keywords

Dopamine; COVIS; Parkinson's disease; positive affect; computational modeling

1 Introduction

Dopamine (DA) is a prominent neuromodulator that is found in many different brain areas. Cools (2006) suggested that prefrontal dopamine levels are related to cognitive stability whereas striatal dopamine levels are related to cognitive plasticity. With such a wide ranging role, almost all cognitive activities should be affected by dopamine levels in the brain. Not surprisingly, factors influencing brain dopamine levels have been shown to affect performance in many behavioral experiments (for a review, see Cools, 2006). For this reason, computational cognitive neuroscience models increasingly include a role for DA in their processing (e.g., Ashby, Alfonso-Reese, Turken, & Waldron, 1998; Ashby & Casale, 2003; Moustafa & Gluck, 2010).

© 2012 Elsevier Ltd. All rights reserved.

For correspondence, Sébastien Hélie, Department of Psychological & Brain Sciences, University of California, Santa Barbara, Santa Barbara, CA 93106-9660, Phone: (805) 284-9474, Fax : (805) 893-4303, sebastien.helie@psych.ucsb.edu.

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

Because DA plays such an important role in cognition, cognitive neuroscientists have worked on identifying experimental manipulations and conditions that can affect brain DA levels. For instance, positive affect (e.g., good mood) is thought to increase the amount of cortical DA (Ashby, Isen, & Turken, 1999). As another example, Parkinson's disease (PD) is caused by the death of DA producing cells in the substantia nigra pars compacta (SNpc) and the ventral tegmental area (VTA), which result in reduced DA levels in the striatum and the prefrontal cortex (Cools, 2006). In this article, we simulate the effects of positive affect and PD on cortical and basal ganglia DA levels in a computational model based on the COVIS theory of categorization (Ashby et al., 1998). The results of three simulations show that DA levels in COVIS modulate performance in a manner that mimics the effect of DA imbalance in humans.

2 The COVIS theory of category learning

COVIS (Ashby et al., 1998) is a neurobiologically detailed theory of category learning that postulates two systems that compete throughout learning – an explicit, hypothesis-testing system that uses logical reasoning and depends on working memory and executive attention, and an implicit system that uses procedural learning. The explicit, hypothesis-testing system of COVIS is thought to mediate *rule-based* category learning. Rule-based category-learning tasks are those in which the category structures can be learned via some explicit reasoning process. Frequently, the rule that maximizes accuracy (i.e., the optimal rule) is easy to describe verbally. In the most common applications, only one stimulus dimension is relevant, and the observer's task is to discover this relevant dimension and then to map the different dimensional values to the relevant categories. The Wisconsin Card Sorting Test (WCST; Heaton, Chelune, Talley, Kay, & Curtiss, 1993) is a well-known rule-based task. More complex rule-based tasks can require attention to multiple stimulus dimensions. For example, any task where the optimal strategy is to apply a logical conjunction or disjunction is rule-based. The key requirement is that the optimal strategy can be discovered by logical reasoning and is easy for humans to describe verbally.

The implicit procedural-learning system of COVIS is hypothesized to mediate *information-integration* category learning. Information-integration tasks are those in which accuracy is maximized only if information from two or more stimulus components (or dimensions) is integrated at some pre-decisional stage. Perceptual integration could take many forms – from treating the stimulus as a Gestalt to computing a weighted linear combination of the dimensional values. Typically, the optimal strategy in information-integration tasks is difficult or impossible to describe verbally. Rule-based strategies can be applied in information-integration tasks, but they generally lead to sub-optimal levels of accuracy because rule-based strategies make separate decisions about each stimulus component, rather than integrating this information.

3 Dopamine imbalances and their effect on rule-based tasks

3.1 Dopamine depletion

The effect of DA depletion in rule-based tasks can be assessed by reviewing the literature on PD. Rule-based tasks demand attention, working memory, and logical reasoning and PD patients display many of the same deficits in these tasks as patients with frontal lobe damage (Owen, Roberts, Hodges, & Robbins, 1993). This section reviews empirical evidence for rule-related deficits in PD patients, with a focus on ineffective use of feedback and proactive interference. Note that other rule-based PD deficits, such as rule-based category learning and perseverative tendencies in the WCST have already been simulated by a COVIS-based model (Helie, Paul, & Ashby, 2011).

Cools, Altamirano, and D'Esposito (2006) asked subjects to predict the outcome of a rule-based gambling task where one stimulus was associated with a reward while the other was associated with a punishment. The stimulus-outcome assignments periodically changed during the task (reversal trials). On reversal trials, subjects received either unexpected positive or unexpected negative feedback. Interestingly, PD patients performed worse in the unexpected negative feedback condition than controls. This suggests that PD patients may have an inability to leverage negative feedback appropriately in this reversal learning paradigm. Similarly, another study using a probabilistic task (Frank, Seeberger, & O'Reilly, 2004) found that PD patients were better at learning from positive than from negative feedback. Together, these results suggest that learning from negative feedback is less effective relative to positive feedback in PD patients. Control (age-matched) subjects in both tasks did not show such differential learning performances with positive and negative feedback.

PD patients also suffer from proactive interference in rule application. In testing PD patients in the Odd-Man-Out (OMO) choice discrimination task (a task where subjects need to pick the odd-man-out in a grouping of three stimuli), Flowers and Robertson (1985) found that PD patients were relatively unimpaired on the first block of trials using one rule (i.e., performance was quite close to controls), but were subsequently impaired in later blocks using either a different rule, or the same original rule. In fact, their performance decrement was only slightly improved when told explicitly what rule to use when selecting the OMO stimulus: subjects never reacquired the same performance level as at the beginning of the test. In a similar task where subjects were required to alternate their response strategy on a trial-by-trial basis, PD patients produced more false alarms than controls, but only for long time intervals between targets (Ravizza & Ivry, 2001). Taken together, these results show that PD patients are sensitive to proactive interference in which early response strategies negatively impact later performance.

3.2 Dopamine elevation

While the studies described above dealt with DA depletion, DA elevations are also thought to have an important impact on rule-based processing. Ashby and his colleagues (1999) reviewed evidence suggesting that cortical dopamine levels are elevated during periods of positive affect. First, dopamine neurons are known to increase their firing following unexpected rewards (e.g., Schultz, Dayan, & Montague, 1997), and the giving of an unexpected reward (e.g., a gift) is a common method of inducing positive affect in test subjects (e.g., Ashby et al., 1999). Second, drugs that mimic the effects of dopamine (i.e., dopamine agonists) or that enhance dopaminergic activity, elevate feelings (e.g., Beatty, 1995). These drugs include morphine and apomorphine (agonists), cocaine (which blocks reuptake), amphetamines (which increase dopamine release), and naturally produced endorphins (e.g., Beatty, 1995; Harte, Eifert, & Smith, 1995). Finally, dopamine antagonists (i.e., neuroleptics) are thought to flatten affect.

Many studies have shown that positive affect improves creative problem solving, facilitates recall of some material, and generally facilitates cognitive flexibility, and Ashby et al. (1999) proposed that these performance improvements were largely due to the elevations in cortical dopamine levels that occur as a result of the improved affect. More recently, Nadler, Rabi, and Minda (2010) studied the effects of positive affect on rule-based and information-integration category learning. Before categorization training, standard methods were used to induce a neutral or positive affect in each subject (i.e., listening to music and watching videos). Results showed that relative to the neutral affect controls, positive affect subjects performed better in rule-based categorization, but not in information-integration categorization.

4 A computational implementation of COVIS

The computational version of COVIS described in this section is an extension of Ashby, Paul, and Maddox (2011) and Hélie et al. (2011). It includes three separate components – namely a model of the hypothesis-testing system, a model of the procedural-learning system, and an algorithm that monitors the output of these two systems and selects a response on each trial. The following subsections describe these components.

4.1 The hypothesis-testing system

The hypothesis-testing system in COVIS selects and tests explicit rules that determine category membership. The simplest rule is one-dimensional. More complex rules are constructed from one-dimensional rules via Boolean algebra (e.g., to produce logical conjunctions, disjunctions, etc.). The neural structures that have been implicated in this process include the prefrontal cortex, anterior cingulate, head of the caudate nucleus, and hippocampus (Ashby et al., 1998, 2005; Hélie, Roeder, & Ashby, 2010). The computational implementation of the COVIS hypothesis-testing system is a hybrid neural network that includes both symbolic and connectionist components. The model's hybrid character arises from its combination of explicit rule selection and switching and its incremental salience-learning component.

To begin, denote the set of all possible explicit rules by $\mathbf{R} = \{R_1, R_2, \dots, R_m\}$. In most applications, the set \mathbf{R} will include all possible one-dimensional rules, and perhaps a variety of plausible conjunction and/or disjunction rules. On each trial, the model selects one of these rules for application by following an algorithm that is described below.

Suppose the stimuli to be categorized vary across trials on r stimulus dimensions. Denote the coordinates of the stimulus on these r dimensions by $\underline{x} = (x_1, x_2, \dots, x_r)$. On trials when the active rule is R_i , a response is selected by computing a discriminant value $h_E(\underline{x})$ and using the following decision rule:

$$\begin{aligned} \text{Respond A on trial } n \text{ if } h_E(\underline{x}) < \varepsilon_E; \\ \text{Respond B if } h_E(\underline{x}) > \varepsilon_E \end{aligned} \quad (1)$$

where ε_E is a normally distributed random variable with mean 0 and variance σ_E^2 . The variance σ_E^2 increases with trial-by-trial variability in the subject's perception of the stimulus and memory of the decision criterion (i.e., perceptual and criterial noise). In the case where R_i is a one-dimensional rule in which the relevant dimension is i , the discriminant function is

$$h_E(\underline{x}) = x_i - C_i \quad (2)$$

where C_i is a constant that plays the role of a decision criterion. Note that this rule is equivalent to deciding whether the stimulus value on dimension i is greater or less than the criterion C_i . The decision bound is the set of all points for which $x_i - C_i = 0$. Note that $|h_E(\underline{x})|$ increases with the distance between the stimulus and this bound.

Suppose rule R_i is used on trial n . Then the rule selection process proceeds as follows. If the response on trial n is correct, then rule R_i is used again on trial $n + 1$ with probability 1. If the response on trial n is incorrect, then the probability of selecting each rule in the set \mathbf{R} for use on trial $n + 1$ is a function of that rule's current weight. The weight associated with each rule is determined by the subject's lifetime history with that rule, the reward history associated with that rule during the current categorization training session, the tendency of

the subject to persevere, and the tendency of the subject to select unusual or creative rules. These factors are all formalized as described next.

Let $Z_k(n)$ denote the salience of rule R_k on trial n . Therefore, $Z_k(0)$ is the initial salience of rule R_k . Rules that subjects have abundant prior experience with have high initial salience, and rules that a subject has rarely used before have low initial salience. In typical applications of COVIS, the initial saliences of all one-dimensional rules are set equal, whereas the initial saliences of conjunctive and disjunctive rules are set much lower. The salience of a rule is adjusted after every trial on which it is used, in a manner that depends on whether or not the rule was successful. For example, if rule R_k is used on trial $n - 1$ and a correct response occurs, then

$$Z_k(n) = Z_k(n - 1) + \Delta_C \quad (3)$$

where Δ_C is some positive constant. If rule R_k is used on trial $n - 1$ and an error occurs, then

$$Z_k(n) = Z_k(n - 1) - \Delta_E \quad (4)$$

where Δ_E is also a positive constant. The numerical value of Δ_C should depend on the perceived gain associated with a correct response and Δ_E should depend on the perceived cost of an error.

The salience of each rule is then adjusted to produce a weight, Y , according to the following rules. For the rule R_i that was active on trial n ,

$$Y_i(n) = Z_i(n) + \gamma \quad (5)$$

where the constant γ is a measure of the tendency of the subject to persevere on the active rule, even though feedback indicates that this rule is incorrect. If γ is small, then switching will be easy, whereas switching is difficult if γ is large. COVIS assumes that switching of executive attention is mediated within the head of the caudate nucleus, and that the parameter γ is inversely related to basal ganglia DA levels.

Choose a rule at random from \mathbf{R} . Call this rule R_j . The weight for this rule is

$$Y_j(n) = Z_j(n) + \mathbf{X} \quad (6)$$

where \mathbf{X} is a random variable that has a Poisson distribution with mean λ . Larger values of λ increase the probability that rule R_j will be selected for the next trial, so λ is called the selection parameter. COVIS assumes that selection is mediated by a cortical network that includes the anterior cingulate and the prefrontal cortex, and that λ increases with cortical DA levels. For any other rule R_k (i.e., $R_k \neq R_i$ or R_j),

$$Y_k(n) = Z_k(n) \quad (7)$$

Finally, rule R_k (for all k) is selected for use on trial $n + 1$ with probability

$$P_{n+1}(R_k) = \frac{Y_k^a(n)}{\sum_{s=1}^m Y_s^a(n)} \quad (8)$$

where a is a parameter that determines the decision stochasticity. When $a < 1$, the decision is noisy and the probability differences are diminished (making the decision probabilities more

uniform). When $a > 1$, the decision tends to become more deterministic. Hence, COVIS assumes that a increases with cortical DA (Ashby & Casale, 2003). This algorithm has a number of attractive properties. First, the more salient the rule, the higher the probability that it will be selected, even after an incorrect trial. Second, after the first trial, feedback is used to adjust the selection probabilities up or down, depending on the success of the rule type. Third, the model has separate selection and switching parameters, reflecting the COVIS assumption that these are separate operations. The random variable \mathbf{X} models the selection operation. The greater the mean of \mathbf{X} (i.e., λ) in Eq. 6, the greater the probability that the selected rule (R_j) will become active. In contrast, the parameter γ from Eq. 5 models switching, because when γ is large, it is unlikely that the system will switch to the selected rule R_j . It is important to note, however, that with both parameters (i.e., λ and γ), optimal performance occurs at intermediate numerical values. For example, note that if λ is too large, some extremely low salience rules will be selected, and if γ is too low then a single incorrect response could cause a subject to switch away from an otherwise successful rule.

4.2 The Procedural System

The current implementation of the procedural system is called the Striatal Pattern Classifier (SPC: Ashby & Waldron, 1999; Ashby, Ennis, & Spiering, 2007). The SPC learns to assign responses to regions of perceptual space. In such models, a decision bound could be defined as the set of all points that separate regions assigned to different responses, but it is important to note that in the SPC, the decision bound has no psychological meaning. As the name suggests, the SPC assumes the key site of learning is at cortical-striatal synapses within the striatum.

The SPC architecture is shown in Figure 1 for an application to a categorization task with two contrasting categories. This is a straightforward three-layer feedforward network with up to 10,000 units in the input layer and two units each in the hidden and output layers. The only modifiable synapses are between the input and hidden layers. The more biologically detailed version of this model proposed in Ashby et al. (2007) included lateral inhibition between striatal units and between cortical units. In the absence of such inhibition, the top motor output layer in Figure 1 represents a conceptual placeholder for the striatum's projection to premotor areas. This layer is not included in the following computational description.

The key structure in the model is the striatum (i.e., the putamen; Ell, Hélie, & Hutchinson, in press; Waldschmidt & Ashby, 2011), which is a major input region of the basal ganglia. In humans and other primates, all of extra-striate cortex projects directly to the striatum and these projections are characterized by massive convergence, with the dendritic field of each medium spiny cell innervated by the axons of approximately 380,000 cortical pyramidal cells (Kincaid, Zheng, & Wilson, 1998). COVIS assumes that, through a procedural-learning process, each striatal unit associates an abstract motor program with a large group of sensory cortical neurons (i.e., all that project strongly to it).

The dendrites of striatal medium spiny neurons are covered in protuberances called spines. These play a critical role in the model because glutamate projections from sensory cortex and DA projections from the SNpc converge (i.e., synapse) on the dendritic spines of the medium spiny cells. COVIS assumes that these synapses are a critical site of procedural learning.

4.2.1 Activation equations—Sensory cortex is modeled as an ordered array of up to 10,000 units, each tuned to a different stimulus. The model assumes that each unit responds maximally when its preferred stimulus is presented, and that its response decreases as a Gaussian function of the distance in stimulus space between the stimulus preferred by that

unit and the presented stimulus. Specifically, when a stimulus is presented, the activation in sensory cortical unit K on trial n is given by

$$I_K(n) = e^{-\frac{d(K, \text{stimulus})^2}{\alpha}} \quad (9)$$

where α is a constant that scales the unit of measurement in stimulus space and $d(K, \text{stimulus})$ is the distance (in stimulus space) between the stimulus preferred by unit K and the presented stimulus (smaller α produces a smaller unit of measurement). Eq. 9 is an example of a radial basis function, a popular method for modeling the receptive fields of sensory units in models of many cognitive tasks.

COVIS assumes that the activation in striatal unit J (within the middle or hidden layer) on trial n , denoted $S_J(n)$, is determined by the weighted sum of activations in all sensory cortical cells that project to it:

$$S_J(n) = \sum_K w_{K,J}(n) I_K(n) + \varepsilon_J \quad (10)$$

where $w_{K,J}(n)$ is the strength of the synapse between cortical unit K and striatal cell J on trial n , $I_K(n)$ is the input from visual cortical unit K on trial n , and ε_J is normally distributed noise (with mean 0 and variance σ_p^2 ; in all the present simulations, $\sigma_p^2 = 0.9$).

In a task with two alternative categories, A and B, the decision rule is:

$$\begin{aligned} &\text{Respond A on trial } n \text{ if } S_A(n) > S_B(n); \\ &\text{Otherwise respond B.} \end{aligned} \quad (11)$$

Hence, smaller σ_p^2 tend to produce more deterministic behaviors. The synaptic strengths $w_{K,J}(n)$ are adjusted up and down from trial-to-trial via reinforcement learning, which is described below.

4.2.2 Learning equations—The three factors thought to be necessary to strengthen cortical-striatal synapses are 1) strong pre-synaptic activation, 2) strong post-synaptic activation, and 3) DA levels above baseline (e.g., see Arbuthnott, Ingham, & Wickens, 2000; Ashby & Hélie, 2011). According to this model, the synapse between a neuron in sensory association cortex and a medium spiny neuron in the striatum is strengthened if the cortical neuron responds strongly to the presented stimulus, the striatal neuron is also strongly activated (i.e., factors 1 and 2 are present) and the subject is rewarded for responding correctly (factor 3). On the other hand, the strength of the synapse will weaken if the subject responds incorrectly (factor 3 is missing), or if the synapse is driven by a cell in sensory cortex that does not produce much activation in the striatum (i.e., factor 2 is missing).

Let $w_{K,J}(n)$ denote the strength of the synapse on trial n between cortical unit K and striatal unit J . COVIS models reinforcement learning as follows:

$$\begin{aligned} w_{K,J}(n+1) = &w_{K,J}(n) \\ &+ \alpha_w I_K(n) [S_J(n) - \theta_{NMDA}]^+ [D(n) - D_{base}]^+ [1 - w_{K,J}(n)] \\ &- \beta_w I_K(n) [S_J(n) - \theta_{NMDA}]^+ [D_{base} - D(n)]^+ w_{K,J}(n) \\ &- \gamma_w I_K(n) [\theta_{NMDA} - S_J(n)]^+ \times [S_J(n) - \theta_{AMPA}]^+ w_{K,J}(n). \end{aligned} \quad (12)$$

The function $[g(n)]^+ = g(n)$ if $g(n) > 0$, and otherwise $g(n) = 0$. The constant D_{base} is the baseline DA level, $D(n)$ is the amount of DA released following feedback on trial n , and α_w , β_w , γ_w , θ_{NMDA} , and θ_{AMPA} are all constants. The first three of these (i.e., α_w , β_w , and γ_w) operate like standard learning rates because they determine the magnitudes of increases and decreases in synaptic strength (in all the simulations herein, $\alpha_w = 0.35$, $\beta_w = 0.45$, and $\gamma_w = 0.15$). The constants θ_{NMDA} and θ_{AMPA} represent the activation thresholds for post-synaptic NMDA and AMPA (more precisely, non-NMDA) glutamate receptors, respectively. The numerical value of $\theta_{NMDA} > \theta_{AMPA}$ because NMDA receptors have a higher threshold for activation than AMPA receptors. This is critical because NMDA receptor activation is required to strengthen corticostriatal synapses (Calabresi, Pisani, Mercuri, & Bernardi, 1992). Note that the values assigned to θ_{NMDA} and θ_{AMPA} are used to discriminate between the postsynaptic activation of the different striatal cells. As such, mid-level values should be selected (because values too high or too low will not allow for such discrimination).

The second line in Eq. 12 describes the conditions under which synapses are strengthened (i.e., striatal activation above the threshold for NMDA receptor activation and DA above baseline) and lines three and four describe conditions that cause the synapse to be weakened. The first possibility (line 3) is that post-synaptic activation is above the NMDA threshold but DA is below baseline (as on an error trial), and the second possibility is that striatal activation is between the AMPA and NMDA thresholds. Note that synaptic strength does not change if post-synaptic activation is below the AMPA threshold.

4.2.3 Dopamine model—The Eq. 12 model of reinforcement learning requires that we specify the amount of DA released on every trial in response to the feedback signal (the $D(n)$ term). The key empirical results are (Schultz et al., 1997): 1) midbrain DA cells fire spontaneously (i.e., tonically), 2) DA release increases above baseline following unexpected reward, and the more unexpected the reward the greater the release, and 3) DA release decreases below baseline following unexpected absence of reward, and the more unexpected the absence, the greater the decrease. One common interpretation of these results is that over a wide range, DA firing is proportional to the reward prediction error (RPE):

$$\text{RPE} = \text{Obtained Reward} - \text{Predicted Reward} \quad (13)$$

A simple model of DA release can be built by specifying how to compute Obtained Reward, Predicted Reward, and exactly how the amount of DA release is related to the RPE. Our solution to these three problems is as follows.

In applications that do not vary the valence of the rewards (e.g., in designs where all correct responses receive the same feedback, as do all errors), the obtained reward R_n on trial n is defined as +1 if correct or reward feedback is received, 0 in the absence of feedback, and -1 if error feedback is received. COVIS computes the predicted reward on trial n from the single-operator learning model (Bush & Mosteller, 1955):

$$P_n = P_{n-1} + .025(R_{n-1} - P_{n-1}) \quad (14)$$

It is well known that when computed in this fashion, P_n converges exponentially to the expected reward value and then fluctuates around this value until reward contingencies change.

Bayer and Glimcher (2005) reported activity in midbrain DA cells as a function of RPE. A simple model that nicely matches their results is:

$$D(n) = \begin{cases} D_{max} & \text{if } RPE > \frac{D_{max} - D_{base}}{D_{slope}} \\ D_{slope} \times RPE + D_{base} & \text{if } -\frac{D_{base}}{D_{slope}} \leq RPE \leq \frac{D_{max} - D_{base}}{D_{slope}} \\ 0 & \text{if } RPE < -\frac{D_{base}}{D_{slope}} \end{cases} \quad (15)$$

where D_{max} , D_{slope} , and D_{base} are constants. This model is illustrated in Figure 2. Note that the baseline DA level is D_{base} (i.e., when the $RPE = 0$) and that DA levels increase linearly with the RPE. In general, higher values of D_{max} allow for a larger increase in DA following unexpected reward, higher values of D_{base} allow for a larger decrease of DA following the unexpected absence of reward, and higher values of D_{slope} increase the effect of RPE on DA release. Thus, increasing the value of any of these constants should improve learning in the procedural system (up to a point).

4.3 Resolving the competition between the hypothesis-testing and procedural systems

Since on any trial the model can make only one response, the final task is to decide which of the two systems will control the observable response. In COVIS, this competition is resolved by combining two factors: the confidence each system has in the accuracy of its response, and how much each system can be trusted. In the case of the hypothesis-testing system, confidence equals the absolute value of the discriminant function $|h_E(n)|$. When $|h_E(n)| = 0$, the stimulus is exactly on the hypothesis-testing system's decision bound, so the model has no confidence in its ability to predict the correct response. When $|h_E(n)|$ is large, the stimulus is far from the bound and confidence is high. In the procedural system, confidence is defined as the absolute value of the difference between the activation values in the two striatal units:

$$|h_P(n)| = |S_A(n) - S_B(n)| \quad (16)$$

The logic of Eq. 16 is similar to that of the hypothesis-testing system: When $|h_P(n)| = 0$, the stimulus is equally activating both striatal units, so the procedural system has no confidence in its ability to predict the correct response, and when $|h_P(n)|$ is large, the evidence strongly favors one response over the other. One problem with this approach is that $|h_E(n)|$ and $|h_P(n)|$ will typically have different upper limits, which makes them difficult to compare. For this reason, these values are normalized to a $[0, 1]$ scale on every trial. This is done by dividing each discriminant value by its maximum possible value.

The amount of trust that is placed in each system is a function of an initial bias toward the hypothesis-testing system, and the previous success history of each system. On trial n , the trust in each system increases with the system weights, $\theta_E(n)$ and $\theta_P(n)$, where it is assumed that $\theta_E(n) + \theta_P(n) = 1$. In typical applications, COVIS assumes that the initial trust in the hypothesis-testing system is much higher than in the procedural system, partly because initially there is no procedural learning to use. A common assumption is that $\theta_E(1) = 0.99$ and $\theta_P(1) = 0.01$. As the experiment progresses, feedback is used to adjust the two system weights up or down depending on the success of the relevant component system. This is done in the following way. If the hypothesis-testing system suggests the correct response on trial n then

$$\theta_E(n+1) = \theta_E(n) + \Delta_{OC} [1 - \theta_E(n)] \quad (17)$$

where Δ_{OC} is a parameter. If instead, the hypothesis-testing system suggests an incorrect response then

$$\theta_E(n+1) = \theta_E(n) - \Delta_{OE} \theta_E(n) \quad (18)$$

where Δ_{OE} is another parameter. The two regulatory terms on the end of Eqs. 17 and 18 restrict $\theta_E(n)$ to the range $0 < \theta_E(n) < 1$. Finally, on every trial, $\theta_P(n+1) = 1 - \theta_E(n+1)$. Thus, Eqs. 17 and 18 also guarantee that $\theta_P(n)$ falls in the range $0 < \theta_P(n) < 1$. The value assigned to Δ_{OC} should be positively related to the model persistence toward hypothesis testing, whereas the value assigned to Δ_{OE} should be positively related to the model willingness to switch to the procedural system.

The last step is to combine confidence and trust. This is done multiplicatively, so the overall system decision rule is: emit the response suggested by the hypothesis-testing system if $|\theta_E(n) - h_E(n)| > |\theta_P(n) - h_P(n)|$; otherwise emit the response suggested by the procedural system.

5 Modeling dopamine imbalances with COVIS

5.1 Parkinson's disease

DA cells in the SNpc and the VTA die in PD, which results in decreased DA levels in the prefrontal cortex and the striatum. In COVIS, DA has a differential effect on the hypothesis-testing and procedural systems. In the hypothesis-testing system, rule selection should improve as levels of DA rise in frontal cortex (up to some optimal level), and rule switching should improve if levels of DA rise in the head of the caudate nucleus. Thus, the selection parameter λ should increase with DA levels in frontal cortex, and the switching parameter γ is assumed to decrease with increased DA levels in the caudate nucleus. In addition, DA in the prefrontal cortex is hypothesized to increase signal-to-noise ratio (Ashby & Casale, 2003). Hence, a in Eq. 4 should increase with DA levels (similar to λ), and σ_E^2 should decrease with more DA (similar to γ). In the procedural system, DA plays a crucial role in learning: it provides the reward signal required for reinforcement learning. A decreased DA baseline or range can affect the ability of the procedural system to learn stimulus-response associations. Hence, decreasing DA levels in the striatum should decrease the values assigned to D_{base} , D_{slope} , and D_{max} .

5.2 Positive affect

According to Ashby et al. (1999), positive affect increases DA levels in frontal cortex, with a much smaller effect in the striatum. This is because the concentration of the dopamine re-uptake molecule DAT is high in the striatum and low in cortex. Thus, dopamine released to the events that induce the positive affect will be cleared quickly from the striatum and slowly from cortex. In COVIS, frontal cortex plays a significant role only in the explicit (hypothesis-testing) system. As argued in Section 5.1, rule selection should improve as levels of DA rise in frontal cortex (up to some optimal level). Thus, the selection parameter λ should increase with cortical DA levels. In addition, DA in the prefrontal cortex increases signal-to-noise ratio (Ashby & Casale, 2003). Hence, a in Eq. 4 should increase with DA levels (similar to λ), and σ_E^2 should decrease.

5.3 Other factors affecting dopamine levels

Many factors are known to affect brain DA levels including age, genetic predisposition, drug-taking history, and neuropsychological patient status (Ashby et al., 1999). For example, brain DA levels are known to decrease by approximately 7% per decade of life due to normal aging, and PD patients are thought to have lost at least 70% of their birth DA levels (Gotham et al., 1988; Price, Filoteo, & Maddox, 2009). Hence, in COVIS, we model an ordinal relationship where $DA(\text{Positive affect; Pos}) > DA(\text{Young adults; YC}) > DA(\text{Old$

adults; OC) DA(PD) (where more DA results in lower γ and σ_E^2 , and higher λ , a , D_{base} , D_{slope} , and D_{max}).¹

Note that D_{base} and D_{max} were calculated to reflect the proportion of DA cells remaining as a function of age and diagnosis (Helie et al., 2011). For instance, in the studies considered here, young adults (YC) are usually undergraduate students in their late teens or early 20s. Hence, they should have approximately 86% of their birth DA levels (assuming they lost 7% of birth DA per decade of life). Typically, these subjects have been modeled with $D_{base} = 0.20$ and $D_{max} = 1.00$ (e.g., Ashby & Crossley, 2011; Ashby et al., 2011). Likewise, age-matched controls (OC) are typically about 70 years old and should thus have 50% of their birth DA level. As such, their D_{base} was set to 0.15 and their D_{max} was set to 0.60. Finally, on average PD patients are predicted to have 30% of their birth DA remaining. Hence, their D_{base} was set to 0.10 and their D_{max} was set to 0.35. Thus, only five DA-related parameters were varied in the simulations (i.e., γ , σ_E^2 , λ , a , and D_{slope}).

6 Simulations

In this section, we test the COVIS model of DA imbalance against data from three well-known tasks, namely rule-based reversal learning, the OMO task, and perceptual categorization. The first two tasks focused on DA depletion and compared the performance of YC and OC with PD patients. In contrast, the last task focused on DA elevations and compared the performance of subjects with neutral and positive affect. The values given to the DA-related parameters in all simulations are shown in Table 1. Note that only these parameters were varied to simulate the different subject populations. In addition to these DA-related parameters, COVIS also requires setting some task-related parameters (which did not vary when modeling the different subject populations). These are shown in Table 2. Note that none of the parameters were optimized; reasonable values were assigned using a rough grid search.

6.1 Ineffective use of negative feedback in Parkinson's disease

The following simulation addresses the ineffective use of negative feedback by PD patients. The key result is that an unexpected reward allows the PD patient to learn and to adjust his or her behavior, but an unexpected punishment is not as effective in eliciting a change in behavior (Cools et al., 2006). In contrast, age-matched controls are as likely to modify their behavior following unexpected reward or punishment.

6.1.1 Experiment—Cools et al. (2006) compared the performance of 10 PD patients with 12 aged-matched controls in a reversal-learning task. On each trial the subject saw two stimuli, one of which was highlighted. The subject's task was to predict whether the highlighted stimulus would lead to a reward or a punishment. The outcome (reward or punishment) was presented after the subject had responded and was non-contingent on the subject's prediction. Hence, the subject needed to generate his or her own internal second-order feedback (e.g., something like "I predicted a reward, the outcome was a reward, therefore my response was correct"). The stimuli were the same on every trial; the only thing that changed was which stimulus was highlighted. At the beginning of each block, one of the stimuli was randomly associated with the reward outcome while the other was associated with the punishment outcome (this assignment was unknown to the subject). After a learning criterion was reached, the previously learned association was reversed. The

¹Note that DA(Pos) = DA(YC) only for cortical DA parameters (i.e., σ_E^2 , λ , and a). For striatal DA parameters (i.e., γ , D_{base} , D_{slope} , and D_{max}), DA(Pos) = DA(YC).

learning criterion was a pseudo-random number of consecutive correct responses that varied between 5 and 9. Each block was composed of 120 trials, and the maximum number of reversals within each block was 14. There were two types of blocks: unexpected reward and unexpected punishment. In the former, the previously punished stimulus was highlighted on the reversal trial and was followed by a reward outcome. In the latter, the previously rewarded stimulus was highlighted on the reversal trial and was followed by a punishment outcome. The dependent measure was the total number of completed stages in each type of block (a stage is a stimulus-reward assignment). The results are shown in Figure 3 (gray bars). As can be seen, the PD patients (right panel) completed more stages in the unexpected reward blocks than in the unexpected punishment blocks. In contrast, the type of block (i.e., unexpected reward or unexpected punishment) did not affect the performance of the aged-matched control subjects (left panel).

6.1.2 Simulation—Two hundred simulations were run for each subject group using the COVIS model described in Section 4. The procedural system received a display-specific representation of the stimuli whereas the hypothesis-testing system received a conceptual representation. Because the stimuli were not perceptually confusable, radial basis functions were not used in the procedural system. Hence, the stimuli presented to the procedural system were 4-dimensional binary vectors. The first two rows coded the first stimulus and the last two rows coded the second stimulus. In both cases, if the stimulus was highlighted a 0 appeared in the first of these two rows and a 1 appeared in the second. If the stimulus was not highlighted the opposite pattern appeared (i.e., 1 in first row, 0 in second). The stimuli presented to the hypothesis-testing system were 2-dimensional binary vectors (one position for each highlighted picture); thus, if stimulus i is highlighted, then row i takes a value of 1. Note that this representation does not take the arrangement of the two stimuli into account because it is a conceptual representation. Each system received a separate copy of the feedback (Helie et al., 2011); however, this experiment required subjects to generate their own second-order feedback. The neuroscience literature suggests that the prefrontal cortex can manipulate highly abstract forms of feedback (Wallis & Kennerley, 2010), whereas the flexibility of the basal ganglia to manipulate feedback is largely unknown (Schultz, Tremblay, & Hollerman, 2000). This distinction was not relevant in previous COVIS simulations because the feedback is generally unambiguous (i.e., both brain areas receive the same feedback), but here the prefrontal cortex and basal ganglia might receive different types of feedback. Specifically, we assumed that the prefrontal cortex receives the self-generated second-order feedback (correct or incorrect dependent on the accuracy of its prediction), whereas the basal ganglia receive the outcome feedback (reward or punishment). Recall that within COVIS, the prefrontal cortex mediates rule selection whereas the basal ganglia mediate rule switching. Hence, when the second-order feedback is negative, but the outcome is positive, the system may try to select a new rule without trying to switch away from the current rule. This makes the likelihood of successfully changing rules less likely than when both the feedback and outcome are congruent (i.e., negative feedback and punishment outcome). Only three free parameters were varied to simulate the data (i.e., γ , λ , and D_{slope} ; see Table 1). The simulation results are shown in Figure 3 (black bars).

6.1.3 Results and discussion—COVIS produced results that were a good match to the human data. As in the Cools et al.'s (2006) results, the performance of the simulated PD patients (right panel) was affected by the valence of the unexpected feedback: unexpected punishment did not lead to a change in response strategy, which ultimately led to fewer stages being completed. In contrast, simulated aged-matched subjects (left panel) were not affected by the valence of the unexpected feedback. These results are explained by the value assigned to the switching parameter γ , which was much larger in simulated PD patients than

in simulated age-matched controls. This parameter assignment is responsible for the perseverative tendencies of simulated PD patients, and resulted in simulated PD patients being unable to disengage from the current stimulus-reward association in unexpected punishment situations. However, in unexpected reward situations, the basal ganglia received positive feedback, which reduced the simulated PD patients' perseverative tendencies and facilitated the switch to a new stimulus-reward association. Aged-matched controls do not suffer from perseverative tendencies (reflected in the value assigned to the switching parameter γ for this group), so the match between the feedback and outcome did not affect their behavior. This good fit was achieved by varying only a subset of the COVIS DA parameters (again, only three parameters were varied). This supports the adequacy of the COVIS DA model with respect to behavioral performance with DA deficits.

6.2 Proactive interference of rule use in Parkinson's disease

The following simulation addresses the proactive interference of rules in PD patients. The key result is that the performance of PD patients is preserved when learning a first rule, but that learning of subsequent rules is impaired (Flower & Robertson, 1985). Control subjects do not suffer from such proactive interference.

6.2.1 Experiment—Flower and Robertson (1985) ran the OMO task to measure the presence of proactive interference in PD patients. The experiment compared the performance of 49 PD patients with 56 aged-matched controls (OC) and 40 younger controls (YC; undergraduate students). The stimuli in the OMO task were two decks of 16 cards. Each card displayed three binary symbols (e.g., circles or triangles) that could each take one of two different sizes (i.e., small or large). The subject's task was to choose a dimension (i.e., symbol or size) and then select the item on each card that was the OMO. For example, consider a card that showed one large circle, one small circle, and one small triangle. If the subject chooses the symbol dimension, the small triangle is the OMO. However, if the subject chooses the size dimension, then the large circle is the OMO. The dimension values for size were always the same (i.e., large, small), but the dimension values for symbols varied within each deck (e.g., circle vs. triangle, rectangle vs. diamond, etc.). In the first block, a deck of cards was selected and the subjects were asked to choose a rule/dimension and consistently apply the same rule for the entire deck of cards. On the second block, the second deck of cards was selected and the subjects were asked to use a different rule for the entire deck of cards (i.e., choose the other dimension). After the second block was completed, the first deck was retrieved and the subjects were asked to apply the first rule again (Block 3). The decks of cards (and rules/dimensions) were alternated until each deck was used four times for a total of 8 blocks. The dependent measure was the number of correct responses in each block (see Figure 4, left panel). As can be seen, the young and old controls exhibited stable performance within each block. However, the PD patients suffered from proactive interference: they made significantly more correct responses in the first block than in any other block.

6.2.2 Simulation—The performance of 200 subjects in each group was simulated with the COVIS model described in Section 4. As in the previous simulation, the procedural system received an object-based representation of the stimuli whereas the hypothesis-testing system received a feature-based representation. Also as in the previous simulation, radial basis functions were not used in the procedural system because the stimuli were not perceptually confusable. Hence, the stimuli presented to the procedural system were 32-dimensional vectors with a 1 in position i to denote stimulus i (for $i = 1, \dots, 32$) and a 0 in every other position. The stimuli presented to the hypothesis-testing system were 6-dimensional binary vectors where every two rows represent a single figure from the stimulus (e.g., rows 1 and 2 encode the size and shape of the first figure in the stimulus). For each consecutive pair of

rows, position i (for $i = 1, 2$) was assigned either a 1 or 0 to indicate the binary value on each dimension (size and shape only took one of two values for each figure). Each system received a separate copy of the feedback. Four free parameters were varied to simulate the data (i.e., σ_E^2 , γ , λ , and D_{stop} ; see Table 1). The simulation results are shown in Figure 4 (right panel).

6.2.3 Results and discussion—As in previous simulations, the model was a good fit to the data. As in the human data, the simulated younger and older controls showed stable performance across all blocks. However, the simulated PD patients suffered proactive interference: this version of the model made significantly more correct responses in the first block than in any other block. This occurred primarily because the PD version of the model had more difficulty both selecting a new rule (lower λ) and switching away from the current rule (higher γ) than the non-PD versions of the model. Note that the first block in this experiment is privileged because it is the only block in which subjects are not already applying a rule when the block begins. Thus, in the first block there is no rule to disengage, so impaired switching and selection have less effect on performance in the first block than on later blocks. Control subjects (both old and young) do not have switching or selection difficulties and therefore show no sign of proactive interference. This good fit to the data was achieved by varying only four DA-related parameters and further supports the adequacy of COVIS as a model of DA depletion.

6.3 The effect of positive affect on cognitive flexibility

The following simulation addresses the effects of positive affect on cognitive flexibility using a category-learning task. The key results are that DA elevations improve explicit rule-based categorization but have no effect on procedural-learning mediated information-integration categorization (Nadler et al., 2010).

6.3.1 Experiment—Nadler and her colleagues (2010) asked 87 university students to listen to music and watch videos that had been shown to induce either a neutral or positive affect in pilot experiments (as measured using a positive affect scale). Following the mood induction, the subjects were tested using a positive affect scale (to ensure that the manipulation worked) and then were trained in either a rule-based or information-integration categorization task for 320 trials. The stimuli in the categorization task were Gabor discs varying in spatial frequency and orientation. In the rule-based condition, the optimal strategy was to ignore the orientation of the discs and apply a rule on frequency (i.e., low frequency stimuli were in category A and high frequency stimuli were in category B). In the information-integration condition, the optimal strategy required subjects to integrate the stimulus-values on the two dimensions at a pre-decisional stage and no simple verbal strategy yielded high accuracy. The dependent measure was categorization accuracy. As seen in Figure 5 (top panels), positive affect subjects performed better than neutral affect subjects in rule-based categorization. However, positive affect did not yield any advantage over neutral affect in information-integration categorization.

6.3.2 Simulation—The performance of 200 subjects in each group was simulated with the COVIS model described in Section 4. Each stimulus was randomly generated using the same distributions (re-scaled to occupy a 100×100 grid) as in Nadler et al. (2010). In the procedural system, the two-dimensional stimuli (frequency, orientation) were used to activate four radial-basis functions located at coordinates (35, 35), (35, 65), (65, 35), and (65, 65) with a common variance of 125 (on both dimensions; the covariance was 0). In the hypothesis-testing system, a rule was selected on each trial and the stimulus value on the rule-dimension was compared with the rule criterion to obtain the discriminant values (as in Eq. 2).² On each trial, a stimulus was selected and processed in each subsystem. Each

system received a separate copy of the feedback. Only two free parameters were varied to simulate the data (i.e., σ_E^2 and λ). The simulation results are shown in Figure 5 (bottom panels).

6.3.3 Results and discussion—As in previous simulations, the model provided a good fit to the data. As in the human data, positive affect improved rule-based category learning but not information-integration category learning. This is because simulated positive affect subjects had more cortical DA (i.e., higher λ and lower σ_E^2), which facilitated rule selection in the hypothesis-testing system and reduced criterial noise. However, cortex plays a limited role in information-integration categorization, which explains the absence of effect in this condition. This good fit to the data was achieved by varying only two cortical DA-related parameters and further supports the adequacy of COVIS as a model of the role of DA in perceptual categorization.

6.4 Parameter sensitivity analysis

As mentioned above, all parameters were selected using a rough grid search approach. Although the parameter values used to simulate the datasets presented were not necessarily optimal, the model's performance is relatively robust to deviations in exact numerical values. Specifically, different numerical values of the parameters within a reasonable range tend to change the performance of the model only slightly. For example, learning rates may change, but not whether the model learns. Thus, we believe that all of the predictions derived in this article follow in a necessary fashion from the general architecture of the model and depend only minimally on our ability to find the optimal set of parameter values.

To verify these observations more formally, we implemented a sensitivity analysis for the simulation of the Nadler and colleagues (2010) data (i.e., Figure 5). The analysis proceeded as follows. For each of the 13 parameters listed in Tables 1 and 2 as well as four parameters described in the text (Equation 10 procedural system noise variance parameter, σ_p^2 , and Equation 12 learning rate parameters α_w , β_w , and γ_w), we successively changed the parameter estimate from the value used to generate the predictions shown in the bottom row of Figure 5 by +10% and -10%. After each change, we simulated the behavior of the model in the same conditions used to generate the bottom row of Figure 5. Next, after each new simulation (and for each condition and categorization task of Nadler et al., 2010), we computed the mean root squared error (MRSE) between the simulated learning curve shown in the bottom row of Figure 5 and the learning curve produced by the new version of the model. Across all parameter adjustments (34), and for every categorization task (RB and II) and condition (positive and neutral affect), the average MRSE was 0.94%, suggesting that the learning performance deviations are very slight with small changes in each parameter. To further corroborate this observation, we re-ran the above analysis modifying each parameter by $\pm 100\%$. The average MRSE across all simulations with this more extreme parameter manipulation was 2.61%. These results suggest that, as long as the parameters are in a reasonable range, no single parameter greatly affects the overall learning performance of the model.

7 General Discussion

This article proposed a formal account of a variety of effects of DA imbalances on cognitive processing. The computational model was based on the COVIS model of categorization (Ashby et al., 1998), and both DA depletion and DA elevation was accounted for by

²The criterion value for each rule was set to the mean value of the stimulus set on the rule dimension.

increasing or decreasing DA-related parameters. The effect of DA depletion was captured by simulating the data of young adults, older adults, and Parkinson's disease patients in a reversal learning task (Cools et al., 2006) and the OMO task (Flower & Robertson, 1985). DA elevation was captured by simulating the effect of positive affect in perceptual categorization (Nadler et al., 2010). The remainder of this article discusses alternative models of DA fluctuations and theoretical implications of the present work.

7.1 Other computational models of dopamine imbalances

Very few computational cognitive neuroscience models of DA imbalances have been proposed. In one of the few, Monchi, Taylor, and Dagher (2000) simulated the effects of DA imbalances to account for the performance of PD patients and schizophrenics in a variety of working memory tasks. Their model includes three basal ganglia-thalamocortical loops: two through the prefrontal cortex (one for spatial information and the other for object information), and one through the anterior cingulate gyrus (for strategy selection). Monchi et al. modeled the reduced DA innervation of the striatum in PD by reducing the connection strengths in the model between cortex and the caudate nucleus, and between the caudate nucleus and the internal segment of the globus pallidus. In contrast, they modeled the effects of schizophrenia by reducing the gain of the units in nucleus accumbens. The model was then used to simulate performance in a delayed response task, a delayed match-to-sample task, and the WCST. In all of these, PD deficits were accounted for by improper encoding of the stimuli in working memory, and schizophrenia deficits were accounted for by a difficulty in selecting the appropriate response strategy.

An alternative model was proposed by Frank (2005) to explain cognitive deficits related to DA depletion in PD patients. This model includes basal ganglia-thalamocortical loops with an emphasis on a more biologically detailed model of the basal ganglia that included both the direct and indirect pathways. In Frank's model, PD is simulated by lesioning SNpc DA cells to reduce the range of DA in the basal ganglia. This reduction in DA's dynamic range reduces activation in the direct pathway (through D1 receptors) and amplifies activation in the indirect pathway (through D2 receptors). In addition, DA plays the role of the reward signal in synaptic plasticity. This model has been used to simulate a probabilistic classification task and a probabilistic reversal learning task. In both tasks, PD deficits were explained by abnormal direct/indirect pathway interactions.

More recently, Moustafa and Gluck (2010, 2011) proposed a new computational model of DA imbalance. Their model is a three-layer feed-forward connectionist network where the input activates the prefrontal cortex, which in turn activates the striatum to produce a response. Similar to Frank (2005), the role of phasic DA is to facilitate synaptic plasticity while the role of tonic DA is to modulate neural activation. However, the Moustafa and Gluck model allows for differential effects of DA in the prefrontal cortex and striatum by varying the slope of the transfer functions and learning rates separately for neurons in these two regions. In Moustafa and Gluck (2011), a model of the hippocampus was added to preprocess the stimuli. PD is simulated by reducing the four DA-related parameters (learning and gain in the PFC and striatum), while schizophrenia is simulated by damaging the hippocampus. The model has been used to simulate instrumental conditioning, probabilistic classification, and probabilistic reversal learning tasks. PD impairments in these tasks were explained by noisy activation and learning while impairments related to schizophrenia were explained by reduced stimulus-stimulus representational learning.

7.2 Theoretical implications

One of the main contributions of the COVIS simulation of DA imbalance is that it brings into focus the different roles of DA in different brain regions. In the Monchi et al. (2000)

and Frank (2005) models, the role of DA is restricted to producing abnormal dynamics in the basal ganglia. Moustafa and Gluck (2010, 2011) were the first to independently simulate the role of DA in the prefrontal cortex and striatum, but the role of DA was the same in both regions: activation gain and learning rate. In COVIS, DA can be independently manipulated in the prefrontal cortex and basal ganglia, but it also has a different role in each region. In the prefrontal cortex, DA facilitates rule selection and increases signal gain (reducing noise). In the basal ganglia, DA facilitates rule switching and synaptic plasticity. These differential roles of DA in the prefrontal cortex and the basal ganglia not only allow for the explanation of a wider range of tasks and phenomena, but also allow for a more fine-grained account of the effects of DA level in each task. For instance, COVIS predicts that in rule-based categorization tasks the primary behavioral effect of DA elevation/depression in the basal ganglia should be to facilitate/impair rule switching (respectively). In information-integration tasks, however, DA fluctuations in the basal ganglia should mostly affect synaptic plasticity (positively or negatively). Previous modeling of DA imbalances did not allow for this level of specificity. In addition, none of the previous models of DA imbalance were used to account for DA elevations, such as those occurring with positive affect. Note that the COVIS dual-process approach successfully modeled reversal learning without a detailed model of the indirect pathway of the basal ganglia.

7.3 Limitation and future work

While COVIS is successful at accounting for many behavioral phenomena observed following various DA imbalances, it cannot yet account for the differential behavioral effects of dopaminergic medication (e.g., Cools et al., 2006; Frank et al., 2004; Gotham et al., 1988). Two of the models reviewed in Section 7.1 have proposed a computational account of the effects of dopaminergic medication (Frank, 2005; Moustafa & Gluck, 2010). We have not made a similar attempt with COVIS for two different reasons. First, Cools et al. (2006) report that different dopaminergic medications may have different behavioral effects. For instance, post hoc analyses suggest that only patients treated with pramipexole were impaired in reversal learning. Most papers where PD patients are tested ON medication report which medications appear in their samples, but do not delineate the ON medication patients according to drug. Hence, it would be difficult to simulate the exact behavioral effects of different PD-related drugs within a particular sample of patients. Second, the issue of medication is further complicated by the observation that dopaminergic treatments have different effects depending on the progression of the disease and this interaction very well could be drug dependent. Thus, while the variable effects of medication on PD performance were not addressed by the current computational model, future work with COVIS could attempt to investigate the differential effects of PD medications when these become reliably reported and controlled in published articles.

Second, it has been hypothesized that schizophrenia may also be characterized by DA imbalances (e.g., Cohen & Servan-Schreiber, 1992). As such, the Monchi et al. (2000) and Mustafa and Gluck (2011) models of DA imbalance have also tried to address schizophrenia. COVIS does not include a detailed model of the hippocampus; however, we would adopt an approach similar to other models that manipulate DA (e.g., Cohen & Servan-Schreiber, 1992; Monchi et al., 2000) to reflect the particular imbalance of DA in schizophrenic patients. For example, DA in the head of the caudate nucleus and the prefrontal cortex could be manipulated. Future work should allow us to determine whether these adjustments to DA parameters in COVIS would produce cognitive deficits similar to those observed in schizophrenia.

Finally, the simulations included in this manuscript were performed using a hybrid symbolic-connectionist implementation of COVIS. More biologically detailed versions of both the procedural and hypothesis-testing systems of COVIS have been implemented and

used to simulate both single-unit recording and behavioral data (Ashby, Ell, Valentin, Casale, 2005; Ashby et al., 2007). However, only behavioral data were simulated in this article. Ashby and Hélie (2011) argued that, while additional biological details can almost always be added to a model, one should only include details that will be tested against empirical data (i.e., the *Simplicity Heuristic*). Hence, because no spike train or single-cell recordings were simulated, the hybrid symbolic-connectionist implementation of COVIS used in this research was appropriate. Even so, future modeling work should focus on neuroscience data where DA levels are directly controlled, and this research should use a more biologically detailed version of COVIS.

Acknowledgments

The authors would like to thank Ruby Nadler for sharing the human data simulated in Section 6.3. This work was supported in part by the U.S. Army Research Office through the Institute for Collaborative Biotechnologies under grant W911NF-07-1-0072, by grant P01NS044393 from the National Institute of Neurological Disorders and Stroke, and by the Intelligence Advanced Research Projects Activity (IARPA) via Department of the Interior (DOI) contract number D10PC20022 (through Lockheed Martin). The U.S. government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of IARPA, DOI, or the U.S. Government.

References

- Arbuthnott GW, Ingham CA, Wickens JR. Dopamine and synaptic plasticity in the neostriatum. *Journal of Anatomy*. 2000; 196:587–596. [PubMed: 10923989]
- Ashby FG, Casale MB. A model of dopamine modulated cortical activation. *Neural Networks*. 2003; 16:973–984. [PubMed: 14692632]
- Ashby FG, Alfonso-Reese LA, Turken AU, Waldron EM. A neuropsychological theory of multiple systems in category learning. *Psychological Review*. 1998; 105:442–481. [PubMed: 9697427]
- Ashby FG, Crossley MJ. A computational model of how cholinergic interneurons protect striatal-dependent learning. *Journal of Cognitive Neuroscience*. 2011; 23:1549–1566. [PubMed: 20521851]
- Ashby FG, Ell SW, Valentin V, Casale MB. FROST: A distributed neurocomputational model of working memory maintenance. *Journal of Cognitive Neuroscience*. 2005; 17:1728–1743. [PubMed: 16269109]
- Ashby FG, Ennis JM, Spiering BJ. A neurobiological theory of automaticity in perceptual categorization. *Psychological Review*. 2007; 114:632–656. [PubMed: 17638499]
- Ashby FG, Hélie S. A tutorial on computational cognitive neuroscience: Modeling the neurodynamics of cognition. *Journal of Mathematical Psychology*. 2011; 55:273–289. [PubMed: 21841845]
- Ashby FG, Isen AM, Turken AU. A neuropsychological theory of positive affect and its influence on cognition. *Psychological Review*. 1999; 106:529–550. [PubMed: 10467897]
- Ashby, FG.; Paul, EJ.; Maddox, WT. COVIS. In: Pothos, EM.; Wills, AJ., editors. *Formal approaches in categorization*. New York: Cambridge University Press; 2011. p. 65–87.
- Ashby FG, Waldron EM. On the nature of implicit categorization. *Psychonomic Bulletin & Review*. 1999; 6:363–378. [PubMed: 12198775]
- Bayer HM, Glimcher PW. Midbrain dopamine neurons encode a quantitative reward prediction error signal. *Neuron*. 2005; 47:129–141. [PubMed: 15996553]
- Beatty, J. *Principles of behavioral neuroscience*. Dubuque, IA: Brown & Benchmark; 1995.
- Bush, RR.; Mosteller, F. *Stochastic models for learning*. New York: Wiley; 1955.
- Calabresi P, Pisani NB, Mercuri NB, Bernardi G. The corticostriatal projection: From synaptic plasticity to dysfunctions of the basal ganglia. *Trends in Neurosciences*. 1992; 19:19–24. [PubMed: 8787136]
- Cohen JD, Servan-Schreiber D. Context, cortex, and dopamine: A connectionist approach to behavior and biology in schizophrenia. *Psychological Review*. 1992; 99:45–77. [PubMed: 1546118]

- Cools R, Altamirano L, D'Esposito M. Reversal learning in Parkinson's disease depends on medication status and outcome valence. *Neuropsychologia*. 2006; 44:1663–1673. [PubMed: 16730032]
- Ell, SW.; Hélie, S.; Hutchinson, S. Putamen: Anatomy, Functions and Role in Disease. Nova Publishers; Contributions of the putamen to cognitive function. (in press).
- Flowers K, Robertson C. The effect of Parkinson's disease on the ability to maintain a mental set. *Journal of Neurology, Neurosurgery & Psychiatry*. 1985; 48:517–529.
- Frank MJ. Dynamic dopamine modulation in the basal ganglia: A neurocomputational account of cognitive deficits in medicated and nonmedicated Parkinsonism. *Journal of Cognitive Neuroscience*. 2005; 17:51–72. [PubMed: 15701239]
- Frank MJ, Seeberger LC, O'reilly RC. By carrot or by stick: cognitive reinforcement learning in parkinsonism. *Science*. 2004; 306:1940–1943. [PubMed: 15528409]
- Gotham AM, Brown RG, Marsden CD. "Frontal" cognitive function in patients with Parkinson's disease "on" and "off" levodopa. *Brain*. 1988; 111:299–321. [PubMed: 3378138]
- Harte JL, Eifert GH, Smith R. The effects of running and meditation on beta-endorphin, corticotropin-releasing hormone and cortisol in plasma, and on mood. *Biological Psychology*. 1995; 40:215–265.
- Heaton, RK.; Chelune, GJ.; Talley, JL.; Kay, GG.; Curtiss, G. Wisconsin Card Sorting Test manual. Odessa, FL: Psychological Assessment Resources, Inc.; 1993.
- Hélie, S.; Paul, EJ.; Ashby, FG. Proceedings of the International Joint Conference on Neural Networks. San Jose, CA: IEEE Press; 2011. Simulating Parkinson's disease patient deficits using a COVIS-based computational model; p. 207-214.
- Hélie S, Roeder JL, Ashby FG. Evidence for cortical automaticity in rule-based categorization. *Journal of Neuroscience*. 2010; 30:14225–14234. [PubMed: 20962243]
- Kincaid AE, Zheng T, Wilson CJ. Connectivity and convergence of single corticostriatal axons. *Journal of Neuroscience*. 1998; 18:4722–4731. [PubMed: 9614246]
- Monchi O, Taylor JG, Dagher A. A neural model of working memory processes in normal subjects, Parkinson's disease and schizophrenia for fMRI design and predictions. *Neural Networks*. 2000; 13:953–973. [PubMed: 11156204]
- Moustafa AA, Gluck MA. A neurocomputational model of dopamine and prefrontal-striatal interactions during multicue category learning by Parkinson patients. *Journal of Cognitive Neuroscience*. 2010; 23:151–167. [PubMed: 20044893]
- Moustafa AA, Gluck MA. Computational cognitive models of prefrontal-striatal-hippocampal interactions in Parkinson's disease and schizophrenia. *Neural Networks*. 2011; 24:575–591. [PubMed: 21411277]
- Nadler RT, Rabi RR, Minda JP. Better mood and better performance: Learning rule-described categories is enhanced by positive moods. *Psychological Science*. 2010; 21:1770–1776. [PubMed: 20974709]
- Owen AM, Roberts AC, Hodges JR, Robbins TW. Contrasting mechanisms of impaired attentional set-shifting in patients with frontal lobe damage or Parkinson's disease. *Brain*. 1993; 116:1159–1175. [PubMed: 8221053]
- Price A, Filoteo JV, Maddox WT. Rule-based category learning in patients with Parkinson's disease. *Neuropsychologia*. 2009; 47:1213–1226. [PubMed: 19428385]
- Ravizza SM, Ivry RB. Comparison of the basal ganglia and cerebellum in shifting attention. *Journal of Cognitive Neuroscience*. 2001; 13:285–297. [PubMed: 11371307]
- Schultz W, Dayan P, Montague PR. A neural substrate of prediction and reward. *Science*. 1997; 275:1593–1599. [PubMed: 9054347]
- Schultz W, Tremblay L, Hollerman JR. Reward processing in primate orbitofrontal cortex and basal ganglia. *Cerebral Cortex*. 2000; 10:272–283. [PubMed: 10731222]
- Waldschmidt JG, Ashby FG. Cortical and striatal contributions to automaticity in information-integration categorization. *Neuroimage*. 2011; 56:1791–1802. [PubMed: 21316475]
- Zakzanis K, Freedman M. A neuropsychological comparison of demented and nondemented patients with Parkinson's disease. *Applied Neuropsychology*. 1999; 6:129–146. [PubMed: 10497689]

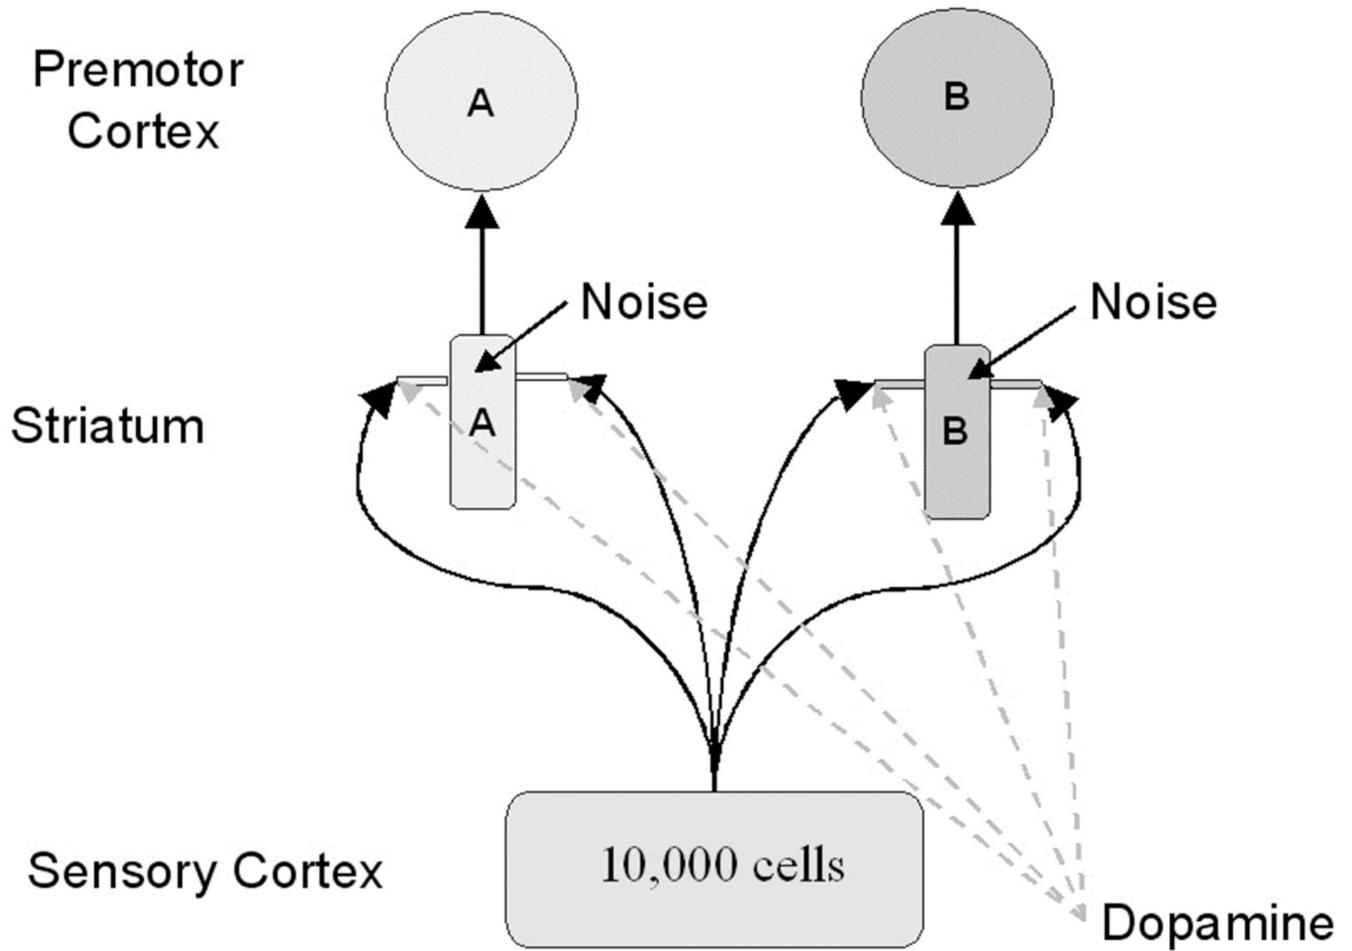


Figure 1.
 A schematic illustrating the architecture of the COVIS procedural system. (From Hélie et al., 2011.)

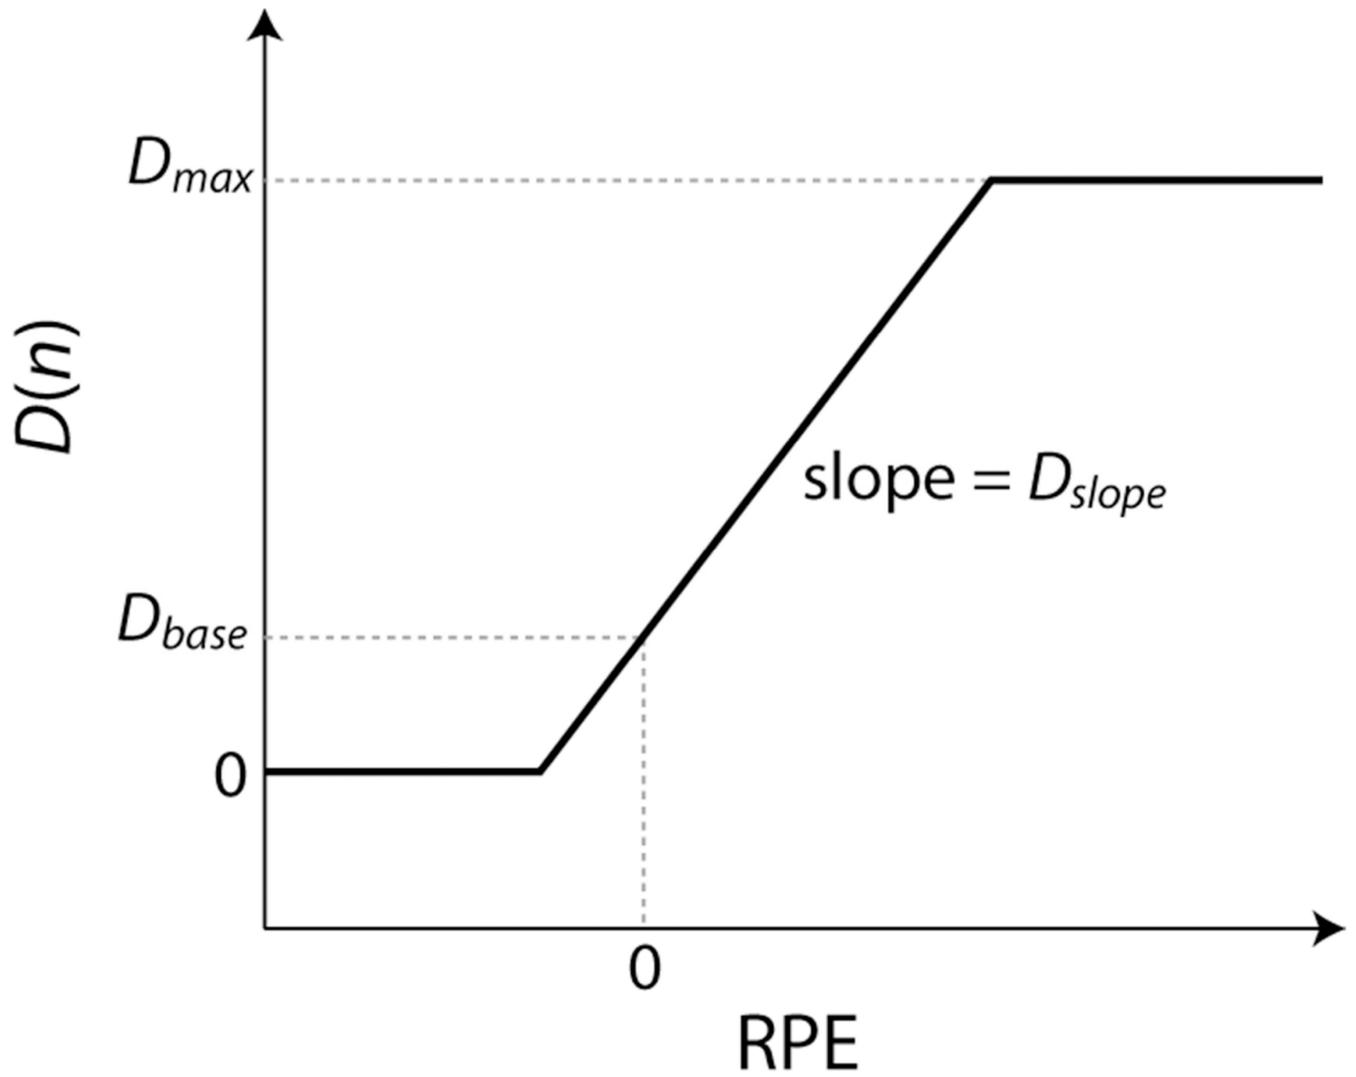


Figure 2. Model used to relate the amount of dopamine (DA) released as a function of the reward prediction error (RPE).

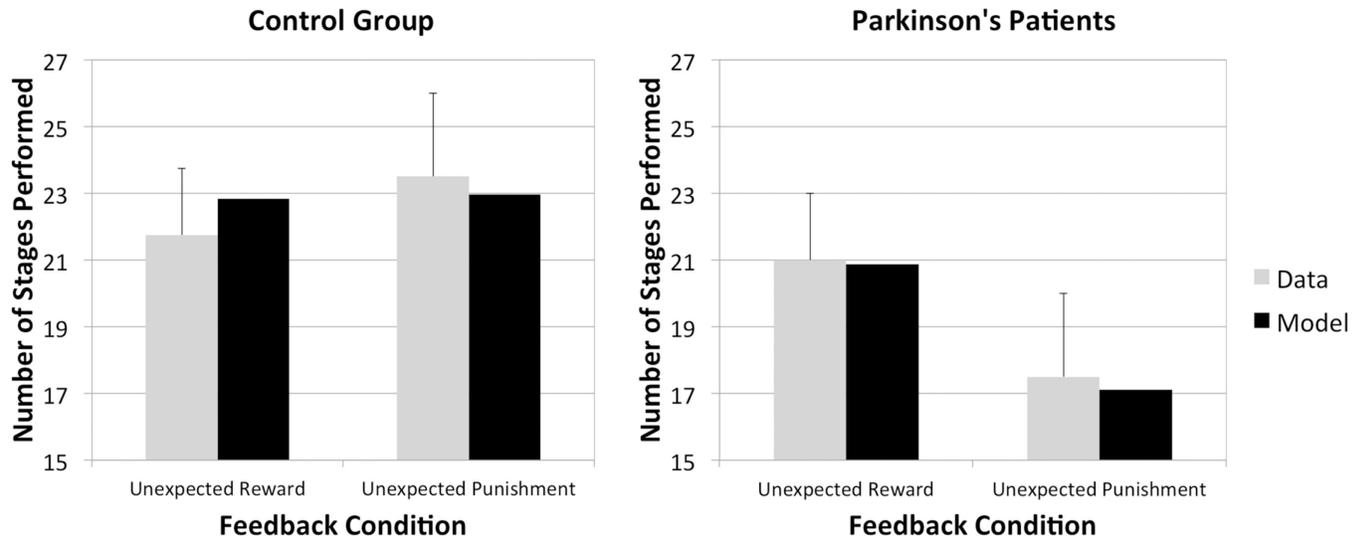


Figure 3. Human (gray bars) and simulation (black bars) data in the reversal learning task of Cools et al. (2006).

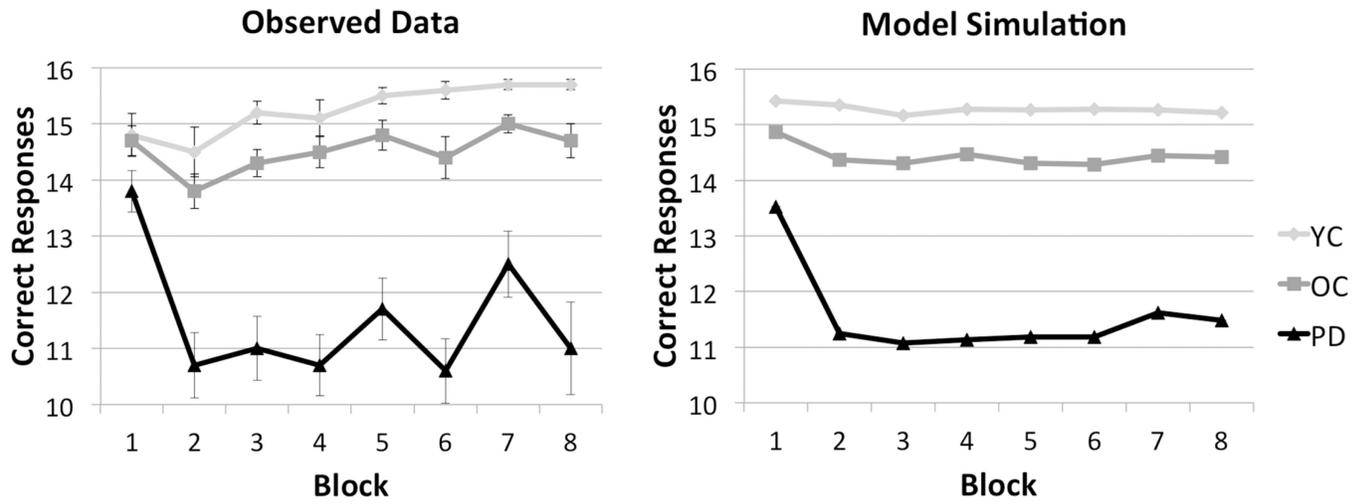


Figure 4. Human (left panel) and simulation (right panel) data in the OMO task of Flowers & Robertson (1985).

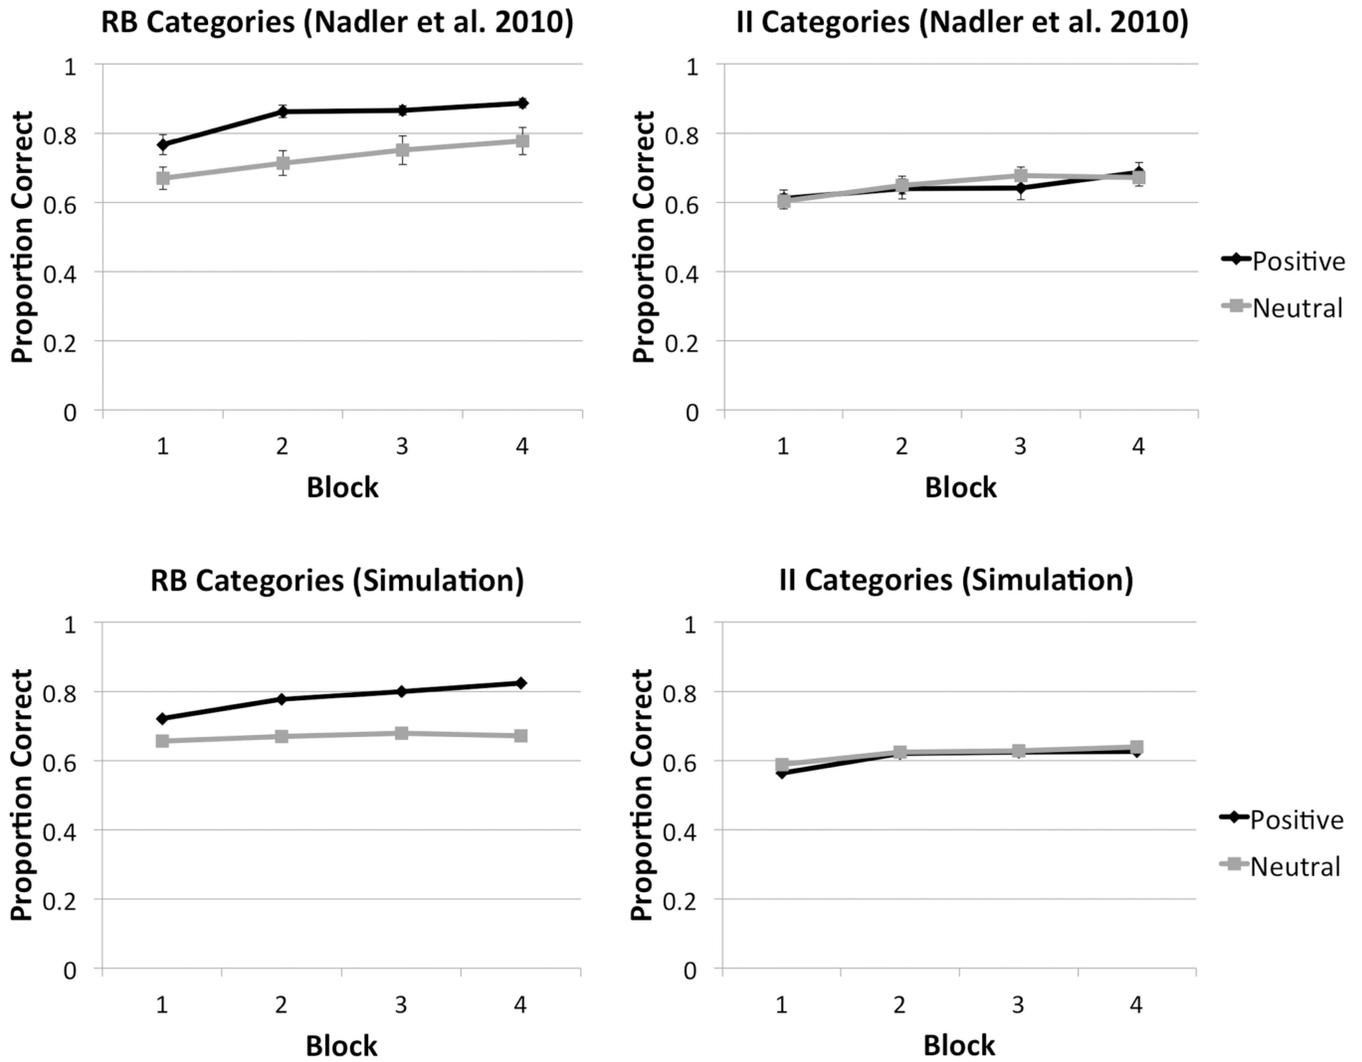


Figure 5. Human and simulation data in the perceptual categorization task of Nadler et al. (2010). The top panels show human results while the bottom panels show the simulation results. RB = Rule-based (verbalizable); II = Information-integration (non-verbalizable).

Table 1

Dopamine-related parameters in COVID

	Cools et al. (2006)	Flowers & Robertson (1985)	Nadler et al. (2010)				
Parameters	OC	PD	YC	OC	PD	Pos	Neu
σ^2_E	0.22	-	0.19	0.23	0.32	0.60	1.50
γ	0.10	1.85	0.25	3.00	25.00	15.00	-
λ	12.1	12.0	9.0	7.0	0.5	7.0	1.0
a	1.1	-	1.0	-	-	1.0	-
D_{slope}	0.80	0.15	0.80	0.25	0.15	0.80	-
D_{max}	0.60	0.35	1.00	0.60	0.35	1.00	-
D_{base}	0.15	0.10	0.20	0.15	0.10	0.20	-

Note. YC = young control; OC = old control; PD = Parkinson's disease patients; Pos = Positive affect; Neu = Neutral affect. Note that D_{max} and D_{base} are not free parameters. They are calculated as a function of the proportion of DA cell lost (see Section 5.3).

Table 2

Task-related parameters in COVIS

Parameters	Cools et al. (2006)	Flowers & Robertson (1985)	Nadler et al. (2010)
Δ_C	0.11	0.09	0.04
Δ_E	0.43	0.05	0.01
θ_{NMDA}	0.057	0.057	0.150
θ_{AMPA}	0.0001	0.0001	0.0010
Δ_{OC}	0.05	0.05	0.04
Δ_{OE}	0.001	0.001	0.010