

A Neurocomputational Model of Dopamine and Prefrontal–Striatal Interactions during Multicue Category Learning by Parkinson Patients

Ahmed A. Moustafa and Mark A. Gluck

Abstract

■ Most existing models of dopamine and learning in Parkinson disease (PD) focus on simulating the role of basal ganglia dopamine in reinforcement learning. Much data argue, however, for a critical role for prefrontal cortex (PFC) dopamine in stimulus selection in attentional learning. Here, we present a new computational model that simulates performance in multicue category learning, such as the “weather prediction” task. The model addresses how PD and dopamine medications affect stimulus selection processes, which mediate reinforcement learning. In this model, PFC dopamine is key for attentional learning, whereas basal ganglia dopamine, consistent with other models, is key for reinforcement and motor learning. The model assumes that competitive dynamics among PFC neurons is the neural mechanism underlying stimulus selection with limited attentional resources,

whereas competitive dynamics among striatal neurons is the neural mechanism underlying action selection. According to our model, PD is associated with decreased phasic and tonic dopamine levels in both PFC and basal ganglia. We assume that dopamine medications increase dopamine levels in both the basal ganglia and PFC, which, in turn, increase tonic dopamine levels but decrease the magnitude of phasic dopamine signaling in these brain structures. Increase of tonic dopamine levels in the simulated PFC enhances attentional shifting performance. The model provides a mechanistic account for several phenomena, including (a) medicated PD patients are more impaired at multicue probabilistic category learning than unmedicated patients and (b) medicated PD patients opt out of reversal when there are alternative and redundant cue dimensions. ■

INTRODUCTION

Parkinson disease (PD) is a neurodegenerative disorder associated with reduced levels of dopamine in the basal ganglia (Jellinger, 1999; Kish, Shannak, & Hornykiewicz, 1988). In addition to basal ganglia dysfunction, several studies showed that mesofrontal dopamine is also affected in PD (Tadaiesky et al., 2008; Prediger et al., 2006; Ashby, Alfonso-Reese, Turken, & Waldron, 1998). Furthermore, some argue that dopaminergic medications (including both the dopamine precursor L-dopa and dopaminergic agonists) increase dopamine levels in prefrontal cortex (PFC; see Silberstein et al., 2005; Kaasinen et al., 2001; Carey, Pinheiro-Carrera, Dai, Tomaz, & Huston, 1995).

Most existing models of learning and cognition in PD focus on simulating the role of basal ganglia dopamine in reinforcement learning. These models assume that mesolimbic dopamine phasic signals projected to the striatum are key for reinforcing motor plans that lead to reward (Guthrie, Myers, & Gluck, 2009; Moustafa & Maida, 2007; Daw, Niv, & Dayan, 2005; Frank, 2005; Suri & Schultz, 1999). Most of these models ignore a potential role for mesofrontal dopamine in behavioral performance. However, existing experimental studies argue for a critical role

for PFC dopamine in stimulus selection processes in attentional learning (Iba & Sawaguchi, 2003). Furthermore, some models of schizophrenia posit that PFC is key for top-down attentional control of motor responding (Amos, 2000; Servan-Schreiber, Bruno, Carter, & Cohen, 1998; Cohen & Servan-Schreiber, 1992). Using features from the models of Kruschke (2003, 2005), Amos (2000), Braver and Cohen (2000), and Cohen and Servan-Schreiber (1992), here, we provide a new computational model that simulates how PD and dopaminergic medications affect performance in multicue category learning tasks.

In multicue category learning tasks, subjects learn to classify multicue patterns into one of two categories based on corrective feedback about their responses. In these tasks, some cues may be more or less diagnostic of category membership than other cues. One example is the “weather prediction” task, in which subjects classify patterns composed of sets of two to four cards (the “cues”) as being predictive of rain versus sunshine (Fera et al., 2005; Shohamy, Myers, Onlaor, & Gluck, 2004; Gluck, Shohamy, & Myers, 2002; Knowlton, Mangels, & Squire, 1996). Another related task is Shohamy, Myers, Hopkins, Sage, and Gluck’s (2009) “slot machine” task, in which subjects are presented with patterns consisting of three cues, each of which can have two values, and subjects learn to categorize these patterns as predictive of payoff in black or white coins.

Performing multicue category learning tasks relies on attentional processes, and thus, the integrity of PFC. Various studies have argued that such multicue category learning relies on the integrity of the fronto-striatal system. For example, Fera et al. (2005) found that PFC and basal ganglia structures are activated during the performance of the “weather prediction” task. In multicue category learning tasks, categorizing patterns based on a subset of cues can lead to suboptimal behavior (but still better than chance). Given that most subjects reach only suboptimal performance in these tasks (Jahanshahi, Wilkinson, Gahir, Dharminda, & Lagnado, 2010; Shohamy et al., 2009), it is possible that subjects pay attention to only a subset of cues in the presented patterns. Multicue category learning tasks (Jahanshahi et al., 2010; Shohamy et al., 2004, 2009) have become regularly used as a means to assess cognitive function in PD patients. Some studies have found that medicated PD patients are more impaired than both unmedicated PD patients and healthy controls at performing multicue category learning tasks (Jahanshahi et al., 2010). Although some researchers have tested PD patients on the “weather prediction” task (Shohamy et al., 2004; Knowlton et al., 1996), Jahanshahi et al. have tested both medicated and unmedicated PD patients on this task. As described later, the model presented here addresses these findings.

Relevant Existing Experimental Studies

Several studies point to a key role for PFC dopamine in attentional processes. For example, in a recent study, Frank, Moustafa, Haughey, Curran, and Hutchison (2007) found that healthy subjects with high levels of PFC dopamine were better at attentional shifting after receiving negative feedback than subject with low levels of PFC dopamine. Several studies also found that dopaminergic medications enhance working memory (WM) performance (Moustafa, Sherman, & Frank, 2008; Lewis, Slabosz, Robbins, Barker, & Owen, 2005). Overall, these studies suggest that increase of dopamine in PFC enhances prefrontal function (Carey et al., 1995), which, in turn, enhances performance in both attentional and WM processes.

Furthermore, PD and dopaminergic medications have been found to impact attentional learning performance. Importantly, dopamine agonists enhance performance in attentional processes. For example, Cools, Barker, Sahakian, and Robbins (2001) found that dopaminergic medications enhance task switching performance in PD patients. In this task, subjects see letters and digits on the screen, and respond to either the digit or the cue based on the color of the screen (e.g., red screen means response to letters, whereas green screen to digits). The model assumes that these effects of dopamine medications are due to an increase of dopamine levels in PFC (see Cools et al., 2001 for further discussion on this).

Dopaminergic medications have also been shown to decrease the occurrence of perseverative errors (Rutledge

et al., 2009; Owen et al., 1993). Owen et al. (1993) found that frontal patients and PD patients who were tested off their dopaminergic medications showed more perseverative errors on attentional shifting tasks than PD patients tested on dopaminergic medication. Similarly, in a recent article from our group in collaboration with Rutledge and colleagues at NYU (Rutledge et al., 2009), we found similar medication effects in a dynamic foraging task, which tests subjects’ ability to adapt to unwarned reversal of reward contingencies. According to our model, dopamine medications enhance attentional shifting, and thus, lead to a decrease in the occurrence of perseverative errors in shifting and reversal tasks. Based on these data, we assume that dopamine medications enhance attentional learning and shifting by increasing mesocortical dopamine levels.

Based on current models and theories, it is not clear how dopamine, basal ganglia, and PFC interact and what their contributions are to the performance of multiple-cue category learning tasks. To the best of our knowledge, there is no existing biologically based computational model that explains the role of both PFC and dopamine in attentional learning in multiple-cue category learning tasks. Furthermore, existing literature does not clarify if and what differential roles PFC and basal ganglia dopamine play in multicue category learning processes. Specifically, here we provide a computational model that simulates performance in various multiple-cue category learning tasks, including the “weather prediction” and “slot machine” tasks. In this model, PFC dopamine is key for stimulus selection in attentional learning, whereas basal ganglia dopamine, consistent with other models, is key for reinforcement and motor learning. Unlike existing models, the basal ganglia in our model is key for mapping attended-to cues to motor responses. Action selection is the process of selecting and executing a motor response from a set of possible actions (also see Prescott, Montes Gonzalez, Gurney, Humphries, & Redgrave, 2006). Similarly, stimulus selection is the mechanism by which one cue comes to control behavior and motor responding. The computational principles underlying stimulus and action selection learning in the model are the same: Phasic dopamine signals are important for stimulus and action selection learning. In addition, competitive inhibitory dynamics—which we simulate using a winner-take-all network—is the computational principle underlying the selection of a motor response as well as focused attention. In other words, the model learns to select and pay attention to only one cue from the presented multicue pattern. The model takes the form of actor–critic architecture (Houk, 1995), in which the critic is key for reward and feedback-based learning and the actor is key for stimulus and action selection learning.

According to the model, PD is associated with decreased phasic and tonic dopamine levels in both PFC and basal ganglia. We assume that dopaminergic medications, commonly used to treat PD patients, increase tonic dopamine levels (beyond those of healthy controls) in

both the basal ganglia and PFC, but decrease the magnitude of phasic dopamine signaling in these brain structures. Dopamine phasic signaling is the difference in firing rate between tonic firing and elevated stimulus-locked firing of dopamine neurons. Thus, an increase in tonic firing of dopamine cells will decrease the magnitude of phasic dopamine signaling (for graphic illustration of these ideas, see Guthrie et al., 2009). We simulate an increase in tonic dopamine by increasing gain value of a sigmoidal activation function, as previously proposed in models of schizophrenia (Cohen & Servan-Schreiber, 1992). (See Methods and Appendix for more information on simulation details.) In our model, an increase in tonic dopamine levels increases activity of postsynaptic cells, as argued by Schultz (2007). An increase of tonic dopamine levels in the simulated PFC enhances attentional shifting performance. Similarly, we simulate a decrease in tonic dopamine levels in a brain structure (as in simulated PD) by decreasing gain value of a sigmoidal activation in the simulated area. The model also shows that lesioning PFC interferes with attentional learning and attentional shifting.

METHODS

Here, we describe the model architecture. The learning algorithm is described in the Appendix.

Model Architecture

The model architecture is shown in Figure 1. It takes the form of an actor–critic architecture, in which the critic is key for reward and feedback-based learning and the actor is key for stimulus and action selection learning. The critic and the actor influence each other in that the critic sends a teaching signal to the actor to strengthen or weaken stimulus and action selection learning (see equations in the Appendix). The critic is not informed about what action the actor has selected, but it is informed about whether the action made had rewarding consequences. The model is trained using the temporal difference (TD) model. We describe details of the TD model in the Appendix.

The model has four modules: input, PFC, motor response, and dopamine module (see Figure 1). The PFC layer is fully connected to the motor response layer (striatum module). Each unit in the input module represents a cue presented to the network. The input and PFC modules have the same number of nodes. The motor module has three nodes, each representing a different motor response. It is important to note here that although there is regional specificity in PFC and the basal ganglia, for the sake of simplicity, we are treating each of these structures in the model as a single module. Input patterns presented to the network activate their corresponding units in the input module. The input module sends topographic projects to the PFC layer (see Figure 2). Here, we use a winner-take-all network to simulate inhibitory

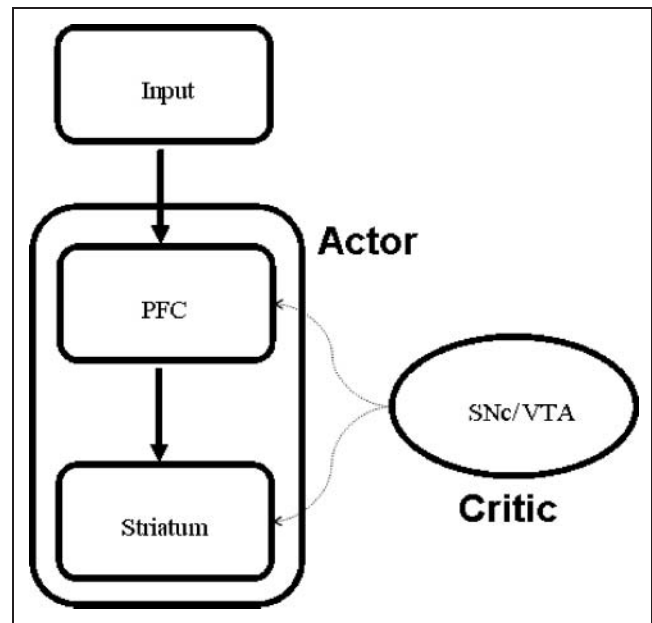
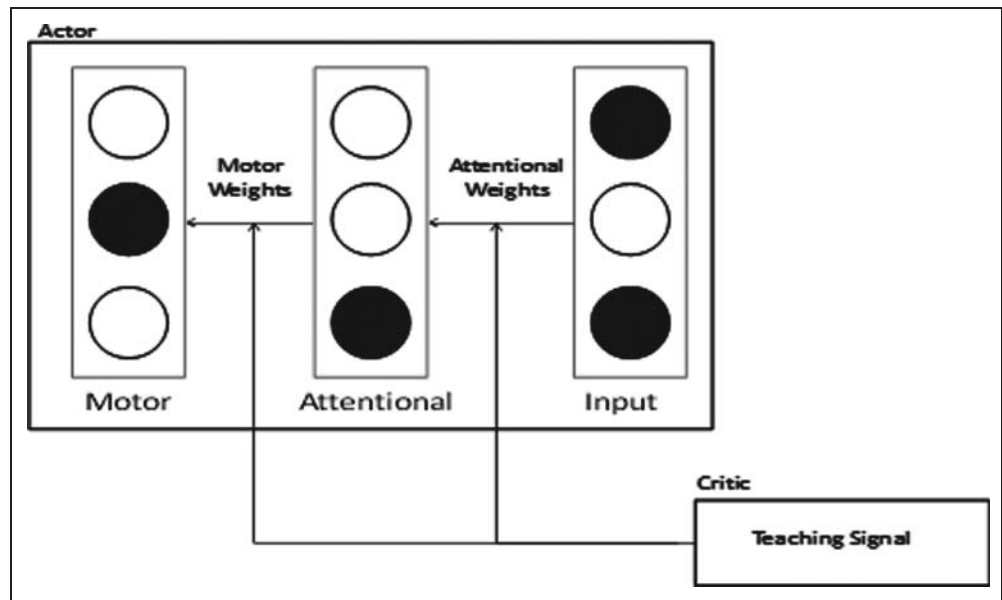


Figure 1. A schematic figure of the model showing relevant brain structures. The model has four modules: input, PFC, motor response, and dopamine module. The critic corresponds to dopamine neurons, whereas the actor corresponds to the prefrontal–striatal system. Learning (i.e., synaptic modification) takes place in both PFC and striatum modules. Learning is modulated by dopamine phasic responses projected from the dopamine module. The input layer sends topographic projections to the PFC layer. The PFC layer is fully connected to the striatum (motor) layer (i.e., every PFC unit is connected to every striatal unit). Activation of a unit in the input layer represents input received from the environment; activation of a unit in the PFC layer represents attended-to stimuli; activation of a unit in the striatum module represents a selected motor response. Dotted lines represent dopaminergic modulatory effects.

connectivity among PFC neurons (see Appendix below for more information on this). At the cognitive level, the winning node represents the attended-to cue. For simplicity, in the current simulations, we allow only one PFC node to be active at each time step. Here, we argue that competitive dynamics among PFC neurons is the brain mechanism underlying limited attentional (as well as WM) processes. Like the Amos (2000) model, we also assume that a negative feedback decreases the activity of most active PFC neurons. As mentioned above, we assume that an increase in tonic dopamine levels increases activity and competition among PFC neurons, which, in turn, enhance selecting different stimuli following negative feedback.

The model has four parameters that are manipulated depending on the simulation of PD and dopaminergic medications. These parameters are two learning rate parameters (one each for the striatal and PFC modules) and two gain parameters (striatal and PFC modules). Learning rate parameters simulate changes in phasic dopamine signaling (see, for example, Shohamy, Myers, Kalanithi, & Gluck, 2008; for experimental support, see Reynolds, Hyland, & Wickens, 2001), whereas gain parameters simulate changes in tonic dopamine levels in the corresponding simulated

Figure 2. A schematic figure of the model showing functions of different modules. The model uses an extended actor–critic architecture in which the critic is responsible for reward-prediction learning and the actor is responsible for action and stimulus selection learning (compare to Figure 1). Like Figure 1, activation of a unit in the input layer represents input received from the environment; activation of a unit in the attentional layer (which is a simulation of PFC) represents attended-to stimuli; activation of a unit in the motor module (which is a simulation of the striatum) represents selected motor response. In the model, the number of attentional weights is equal to the number of input nodes—each input node has a corresponding attentional weight. In this example, a filled node means it is active (and an unfilled node means it is inactive). In this example, Cue 1 and Cue 3 are presented to the model (see activation of nodes in the input module). The model is paying attention to Dimension 3 (see attentional module) and selects motor response, a motor response represented by the middle node in the motor module. Teaching signal corresponds to dopamine phasic signals coming from the ventral tegmental area and substantia nigra pars compacta.



brain structure (Servan-Schreiber, Cohen, & Steingard, 1996; Cohen & Servan-Schreiber, 1992; see Appendix for a description of all parameters). We simulate PD by decreasing learning rate and gain values in the basal ganglia and PFC. We simulate the effects of dopaminergic medication by increasing gain values while concurrently decreasing learning rate values, beyond those used for healthy participants (Table 1).

The simulated striatum in the proposed model learns to map input stimuli to responses (for similar ideas, see Guthrie et al., 2009; Suri & Schultz, 1999). Like the PFC module, we use a winner-take-all network to simulate inhibitory connectivity among simulated striatal neurons. At the cognitive level, the winning node represents the selected motor response (for similar ideas, see Guthrie et al., 2009; Suri & Schultz, 1999). Unlike most existing basal ganglia models (Ashby, Ell, Valentin, & Casale, 2005; Frank, 2005; Suri & Schultz, 1999; Houk, 1995), the basal ganglia, in our model, learns to map representations of *attended-to* stimuli to motor responses. These mechanisms, as discussed below, can explain performance in various multiple-cue category learning tasks.

Based on experimental findings (Silberstein et al., 2005; Kaasinen et al., 2001; Carey et al., 1995), it is likely that dopaminergic medications increase dopamine levels in PFC. Specifically, we simulate an increase in PFC tonic dopamine levels by increasing the gain value of the sigmoidal activation function, as previously proposed by various computational models (Amos, 2000; Cohen & Servan-Schreiber, 1992). This, in turn, increases the signal-to-noise ratio in PFC neurons, which has important implications for attentional learning processes: Increasing the gain value

causes the sigmoidal function to decrease the difference between inputs. For example, the difference in the activation level given input of 0.7 versus 0.8 decreases as the gain value increases; that is, for a large gain value, the sigmoidal activation function de-emphasizes differences in input values (within a certain range). In other words, increasing the gain value increases the competitive dynamics in PFC and the likelihood to shift to a new dimension. This has implications for learning multiple-cue category learning tasks, as described below.

RESULTS

We now describe how the model and theories proposed above can simulate performance in various category learning tasks. We first describe simulation results of instrumental conditioning, “weather prediction,” and some other multiple-cue category learning tasks (Jahanshahi et al., 2010; Shohamy et al., 2009). All simulation results presented here are based on averages across 50 runs of the model (see Table 1 for parameter values chosen for simulation results presented below).

Table 1. Parameter Values Used in the Simulations

<i>Parameter</i>	LR_{BG}	G_{BG}	LR_{PFC}	G_{PFC}
HC	0.13	1	0.06	1
PD off	0.09	0.06	0.032	0.06
PD on	0.06	1.9	0.01	1.9

LR = learning rate; G = gain; BG = basal ganglia; PFC = prefrontal cortex. See Appendix for description of all parameters.

We note that the learning rate is larger for the basal ganglia than for the PFC module to allow for faster motor than attentional learning and, most importantly, to allow the model to explore different responses after a negative feedback before shifting attention to different cues.

Instrumental Conditioning (Stimulus–Response Learning)

Instrumental conditioning is a simple task in which subjects learn to associate different stimuli with different responses, based on corrective feedback. In this task, on each trial, a different cue (A or B) is presented at Time Step 1. The model is trained to associate different cues with different responses (make response R1 if A is presented and R2 if B is presented). Feedback is delivered at Time Step 2: Reward presented is 1 if the model made the correct response and is 0 otherwise. Many experimental studies have shown that PD patients are impaired at instrumental conditioning, or stimulus–response (S–R) learning, tasks (Filoteo, Maddox, Ing, & Song, 2007; Shohamy, Myers, Gehman, Sage, & Gluck, 2006; Czernecki et al., 2002). Czernecki et al. (2002) found that both medicated and unmedicated PD patients were impaired on this task. Similarly, Shohamy et al. (2006) found medicated PD patients to be more impaired than unmedicated patients at feedback-based associative learning. In our simulations, we also found both simulated medicated and unmedicated PD patients to be more impaired at stimulus–reward learning than healthy controls (see average correct performance in each block in Figure 3).

We now describe the performance of the critic and the pattern of changes in the TD error signal during the course of learning (see Appendix for definition of TD error). Figure 4B shows the TD error for the successive time steps

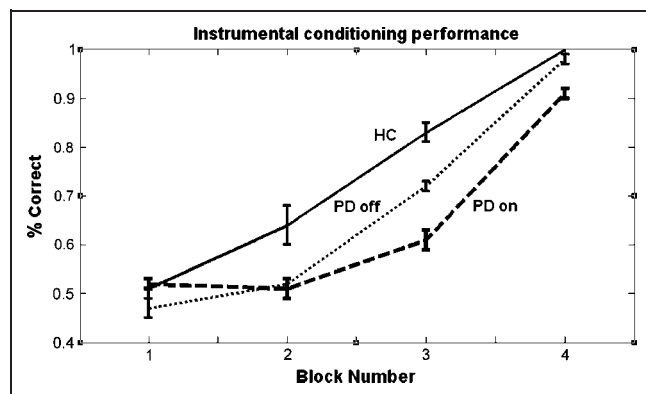


Figure 3. Simulation results of instrumental conditioning (or S–R) performance. The data from all simulation runs were divided into four blocks of 25 trials each, and we show an average correct performance in each block (same applies to all experiments presented below). In the model, both medicated and unmedicated PD patients were more impaired at S–R learning than healthy controls. HC = healthy controls; PD off = Parkinson patients off medications; PD on = Parkinson patients on medications.

within one simulation run. Figure 4D shows the TD error when the model’s response was correct. On each of these trials, the model received the primary reward whose value was 1. As learning proceeded, the TD error signal moved earlier in time to the time step of CS presentation. This is in accordance with the observed phasic response of dopamine neurons (Schultz, Dayan, & Montague, 1997; Figure 4A and C). In subsequent simulation studies, the characteristic of TD error signal is the same as shown here, so we will not present them again.

“Weather Prediction” Task (Jahanshahi et al., 2010)

In the “weather prediction” task, subjects are presented with patterns of cards, each of which is composed of one to three cards taken from a set of four cards; subjects are asked to classify each pattern of cards into one of two categories (Fera et al., 2005; Shohamy et al., 2004; Gluck et al., 2002; Knowlton et al., 1996). On each trial, subjects see one of these patterns, and are asked to predict whether there will be good or bad weather the next day (sun or rain). In this task, each cue (card) is a partially accurate predictor of the outcome category (Fera et al., 2005; Shohamy et al., 2004; Gluck et al., 2002; Knowlton et al., 1996). Jahanshahi et al. (2010) found that medicated PD patients are more impaired at this task than unmedicated patients, who were, in turn, more impaired than healthy controls (see Figure 5A). In agreement with these findings, we found similar qualitative results: Simulated medicated PD patients were more impaired than unmedicated PD patients, who were more impaired than controls (see Figure 5B). In addition, the model provides an account for suboptimal behavior of subjects (including healthy controls) in this task. We argue that suboptimal behavior in this task is due to the model paying attention to one cue of the presented multicue pattern. In this case, the model learns to select and pay attention to the most advantageous card, and categorizes patterns according to this rule. This is also in agreement with experimental finding that human subjects often appear to follow single-cue strategies (i.e., responding on the basis of the presence or absence of a single cue, disregarding all other cues) to solve the “weather prediction” task (Gluck et al., 2002) (see model limitations and future directions below for why this is not always so). Furthermore, a lower learning rate in medicated PD patients than in unmedicated PD patients (see Table 1) explains why dopamine medications impair learning, as reported in Jahanshahi et al. In the model, dopamine medications increase tonic dopamine levels, which, in turn, compromise dopamine phasic signaling.

Multiple-cue Category Learning “Slot Machine” Task (Shohamy et al., 2009)

This task has two phases: acquisition and reversal. In the acquisition phase, subjects are presented with one of eight

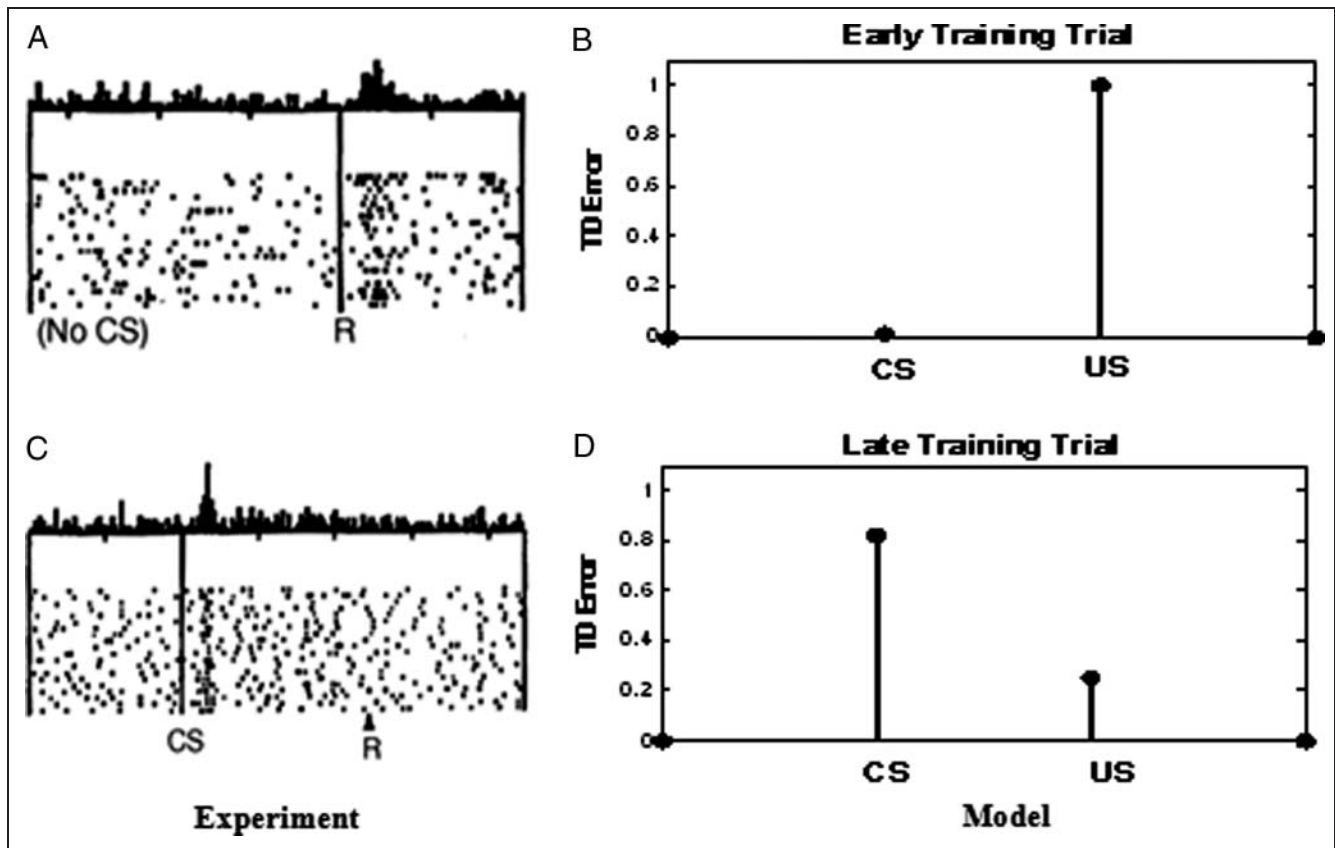


Figure 4. The time shifting of dopamine phasic signal and TD error during instrumental conditioning task learning. (B and D) The TD error captures main characteristics of dopamine phasic responses (A and C) to conditioned and unconditioned stimuli. (A and C) Figures are adapted from Schultz et al. (1997). In subsequent simulation studies, the time shifting of the TD error is very similar to the figures here, so we will not show them again. R = reward; US = unconditioned stimulus; CS = conditioned stimulus.

different patterns; each pattern is composed of three different cues that can each have one of two values (see Table 2). Subjects learn to categorize each pattern based on feedback provided after subjects select a category. Each pattern belongs to either Category A or B. Optimal performance in this task depends on subjects responding to information about all cues presented in each pattern. However, responding based only on one cue from the pre-

sented pattern will lead to suboptimal (and above chance) performance. Upon finishing the acquisition phase, subjects perform a reversal phase, in which the category membership is reversed. Shohamy et al. (2009) found that medicated PD patients were not significantly more impaired than healthy controls in the acquisition and reversal phases (Figure 6A and B). Like Shohamy et al., our simulation results showed that medicated patients' performance

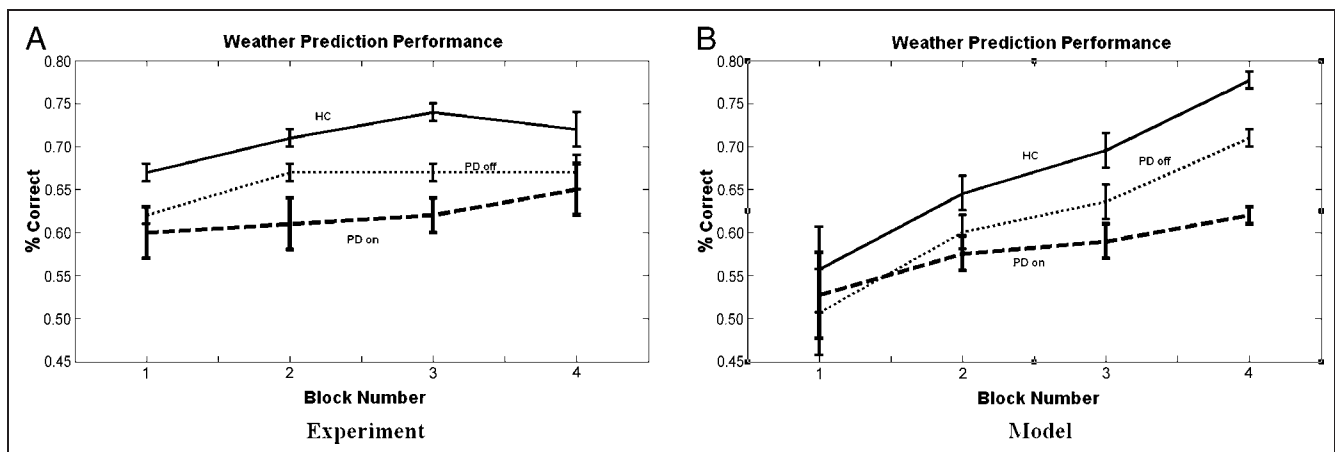


Figure 5. Weather prediction task performance. Each block here has 50 trials. (A) Adapted from Jahanshahi et al. (2010). (B) Simulation results are qualitatively similar to the results of Jahanshahi et al.

Table 2. Description of the “Slot Machine” Task Used by Shohamy et al. (2009) and Other Variants of This Task

<i>Cue 1</i>	<i>Cue 2</i>	<i>Cue 3</i>	<i>Slot Machine Task (Phase 1: Acquisition)</i>	<i>Slot Machine Task (Phase 2: Reversal)</i>	<i>Forced-cue and Shifting (Phase 1: Acquisition)</i>	<i>Forced-cue Version (Phase 2: Reversal)</i>	<i>Shifting Version (Phase 2: Shifting)</i>
1	1	1	A	B	A	B	A
1	1	0	A	B	A	B	A
1	0	1	A	B	A	B	B
1	0	0	B	A	A	B	B
0	1	1	A	B	B	A	A
0	1	0	B	A	B	A	A
0	0	1	B	A	B	A	B
0	0	0	B	A	B	A	B

0 and 1 represent different cues. A pattern consists of three cues presented together. Each row in the first three columns on the left represents a pattern. For example, Pattern 110 represents “candle–fish–plane,” whereas Pattern 101 represents “candle–butterfly–boat.” A and B represent different categories. The slot machine task used by Shohamy et al. (2009) has two phases: acquisition and reversal (see forth and fifth columns from the left). In the reversal phase, we simply change category membership for each pattern (e.g., if A was correct during acquisition, B is correct during reversal). Here, we also propose two task variants of the slot machine task: forced-cue and shifting tasks. Both tasks have the same acquisition phase (see third column from the right). Here, paying attention to one cue is key for a good performance (example here is Cue 1). The second phase of forced-cue task is a reversal phase (see second column from the right). In the second phase of the shifting task, paying attention to a different cue (in this example, it is Cue 2 instead of Cue 1) will lead to a correct performance.

is slightly lower than controls in both the acquisition and reversal phases (Figure 6C and D). Shohamy et al. did not test unmedicated PD patients on this task. The model predicts that unmedicated PD patients will be slightly slower to learn the acquisition phase but perform similarly to healthy controls during the reversal phase. Future experimental work should confirm or refute this prediction. According to the model, because of inhibitory dynamics among PFC neurons, the model can pay attention to only one cue at each time step. Accordingly, the model learns to map the attended-to cue to its corresponding category. This provides a mechanistic account for suboptimal—and above chance—behavior in medicated PD and healthy control subjects, as found in the Shohamy et al. study. An increase of dopamine levels in PFC (which increases activity of PFC neurons), along with receiving negative feedback in the reversal phase, causes medicated PD patients to be more likely to pay attention to a different cue, and thus, opt out of reversal, as reported in the Shohamy et al. study. In other words, an increase in the gain value of sigmoidal activation function of the PFC module increases the competition among PFC nodes, and thus, any of the cues is likely to win the competition (see Discussion for more details).

Furthermore, Shohamy et al. also found that most subjects appeared to follow a one-cue strategy in categorizing the patterns; that is, they categorized the patterns guided by only one of the cues in the presented pattern. In the reversal phase, most healthy control subjects reversed along the same cue they have chosen during the acquisition phase, whereas approximately 65% of medicated PD patients learned the new association by shifting to a one-cue strategy based on a different cue, therefore opting out of reversal. We analyzed changes in attentional weight

values (weights connecting the input module to the PFC module) during task learning (Figure 7). Simulation results showed that all groups paid attention to only one cue during the acquisition and reversal phases. Furthermore, simulation results show that, in most runs, medicated PD patients paid attention to a different cue during the reversal phase. This is explained by attentional weight change during reversal performance (see Figure 7F).

We also conducted strategy analysis on model output for each run. Here, we assume that each run of the model corresponds to a different subject performing the task (for more details on strategy analysis conducted on model output, see Appendix and also Gluck et al., 2002). Simulation results also found that simulated medicated PD randomly chooses to pay attention to one of the three cues during reversal, making it likely to reverse along the same dimension which happens in about third of the simulation runs (Figure 8B). Interestingly, we found that the model behavior here is in agreement with the existing experimental data of the Shohamy et al. study (see Figure 8A). We also predict that similar to healthy controls, unmedicated PD patients will reverse along the same cue chosen during the acquisition phase (Figure 8C).

As we have mentioned earlier, the key feature of this model is simulating the role of PFC dopamine in behavioral performance. Here we study the effects of lesioning PFC on performing the “slot machine” task. We simulated lesions of a brain area by adding noise to the activation levels of the simulated area, as previously done in other computational models (Joanisse & Seidenberg, 1999; Olson & Humphreys, 1999). In the “slot machine” task (Shohamy et al., 2009), the PFC-lesioned model was slower at performing the acquisition and reversal phases; the lesioned model also

showed more perseverative errors than the intact model. This is because adding noise interfered with quickly finding and responding to one cue during early training trials. In other words, adding noise to PFC activity makes strengthening weights in the PFC module in early training trials somewhat ineffective, and thus, slows down learning.

Variants of the “Slot Machine” Task: Forced-cue and Shifting

Here, we simulate performance in two different variants of the original “slot machine” task (Shohamy et al., 2009): forced-cue and shifting tasks (see Table 2 for task description). The acquisition phase is the same in both tasks but is different from the acquisition phase used in the origi-

nal “slot machine” task (see Table 2). In the acquisition phase of both tasks, only one of the three cues is diagnostic of category membership (with 100% probability), whereas the other two are less reliably predictive of category membership (50% probability). This is different from the original task because, in the original task, the configuration of all three cues was predictive of category membership, and each individual cue predicted category membership with only 80% probability. In the reversal phase of the forced-cue task, that same cue is still diagnostic (although the valence is reversed), and thus, subjects cannot opt out of reversal by shifting to another equally diagnostic cue (unlike the original task). In the second phase of the shifting task, the previously diagnostic cue becomes irrelevant, and subjects must shift—and pay attention to a

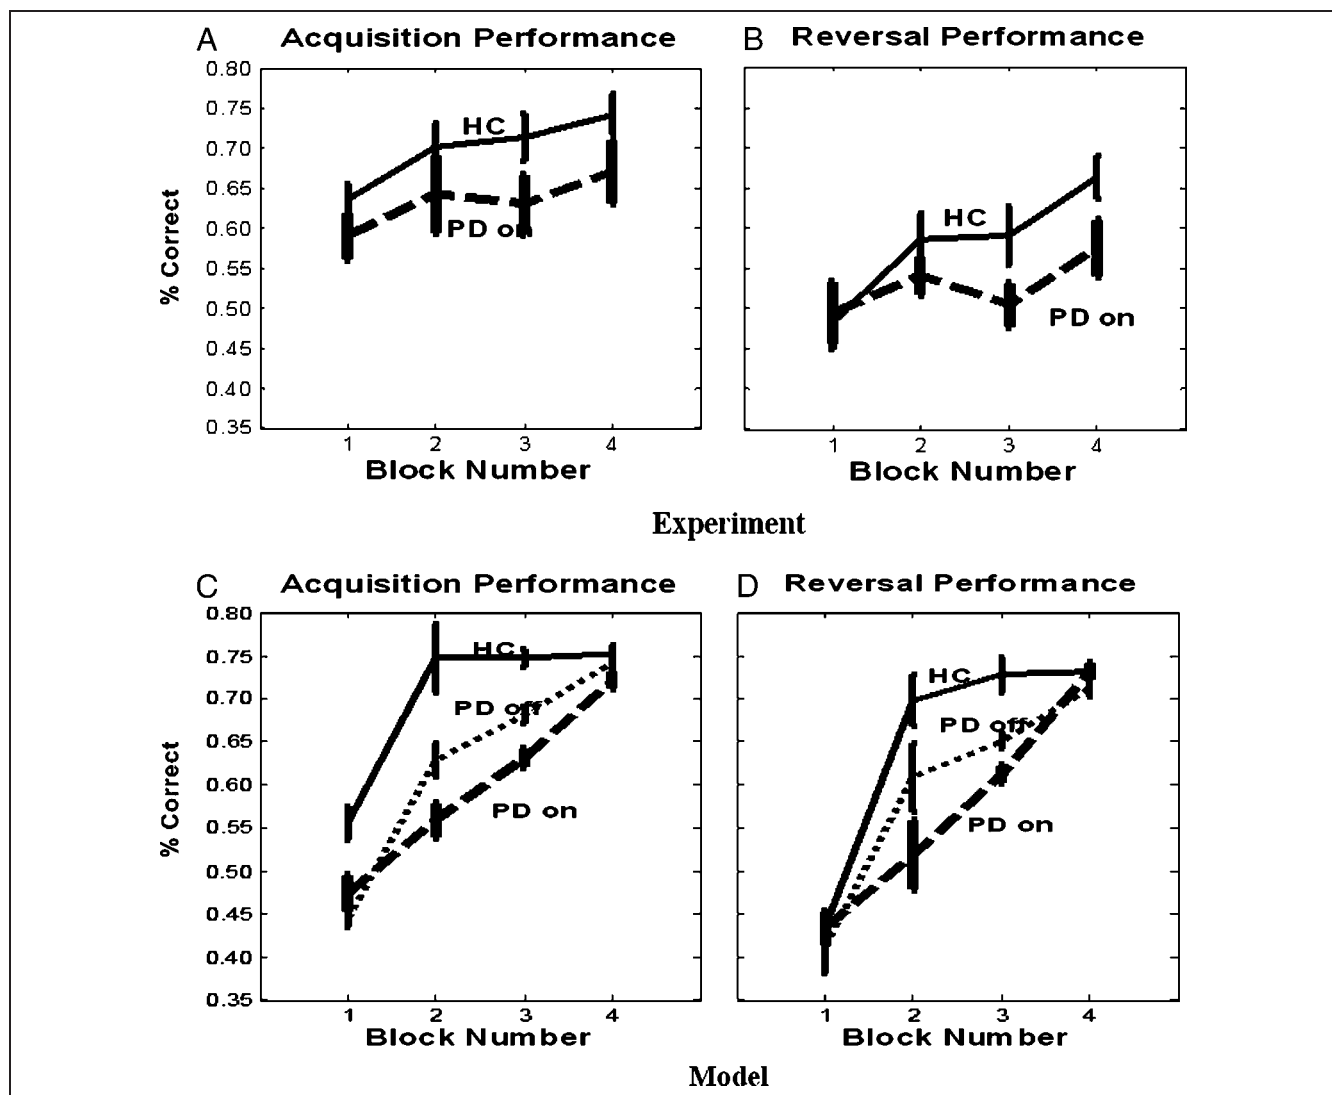


Figure 6. Multicue category learning performance (Shohamy et al., 2009). The data from all simulation runs were divided into four blocks of 25 trials each. (A and B) Experimental results of the acquisition and reversal phases of the multicue category task (Shohamy et al., 2009). Shohamy et al. found medicated PD patients to be numerically worse at both the acquisition and reversal phases of the task. They have not tested unmedicated patients on the task. Figure is adapted from Shohamy et al. (2009). (C and D) Simulations results of the study of Shohamy et al. Simulation results are qualitatively similar to the results of Shohamy et al. The model predicts that unmedicated PD patients would be slightly better at performing this task than medicated patients.

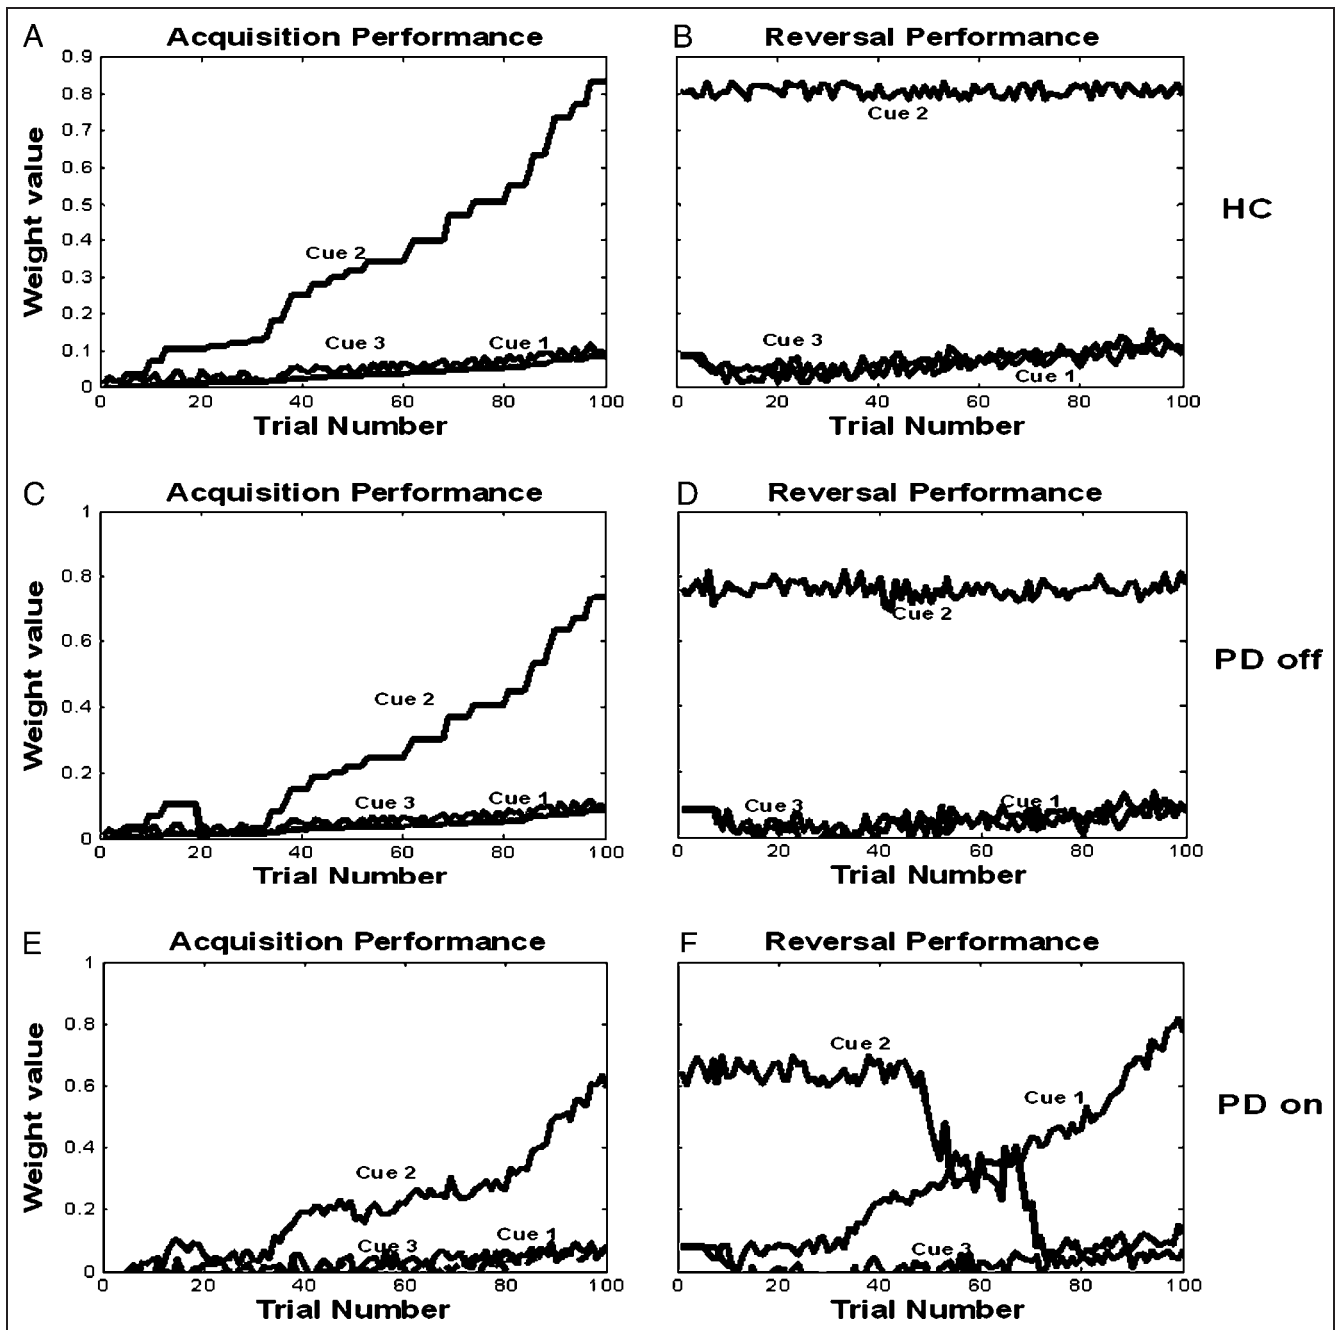


Figure 7. Attentional weight change during the performance of the “slot machine” task in simulated healthy controls (top), unmedicated PD patients (middle), and medicated PD patients (bottom). (A, C, E) Change in attentional weight values during performance in the acquisition phase of the multicue category task (Shohamy et al., 2008). (B, D, F) Change in attentional weight values during performance in the reversal phase of the task. These are data from a typical run in all groups. In all groups, the model selected a random cue and categorized patterns based on that rule. The example shown here is when the model selected Cue 2. In reversal, simulated healthy controls (HC) and PD off patients reversed along the same cue (Cue 2) selected during acquisition. However, in the example shown here, simulated medicated PD (PD on) chose Cue 1 during reversal, therefore opting out of reversal. This is because of increased tonic DA levels in PFC that increase competition among PFC nodes, and thus, also increase the likelihood of shifting attention after receiving negative feedback.

different cue in order to be able to respond correctly (see Table 2).

Unlike the original “slot machine” task, modeling results showed that simulated medicated PD patients are more impaired at reversal learning in the forced-cue task than simulated unmedicated PD patients (see Figure 9). This is due to increased attentional shifting performance in simulated

medicated PD patients, which impairs performance in the reversal phase of this task because the simulated medicated PD patients shift their attention away from the relevant stimuli. Correct performance in this task requires subjects to reverse responses along the same cue chosen in the acquisition phase (see Table 2). In the shifting task, contrary to our expectations, we found that simulated medicated PD

patients perform better than controls and unmedicated patients in the first block of the shifting phase, although the other groups quickly caught up, and controls eventually surpassed both PD groups due to the low learning rates in the PD groups. This is a novel prediction of the computational model, which remains to be tested empirically.

DISCUSSION

Here, we presented a simple model that simulates a potential role for the mesofrontal and nigrostriatal dopamine in

stimulus and action selection in multicue category learning tasks. The model assumes that basal ganglia (nigrostriatal) dopamine is key for motor and S–R learning, whereas PFC (mesocortical) dopamine is key for stimulus selection learning. The model provides a qualitative fit to some of the existing behavioral data (but see below for model limitations and potential future extension to the model).

In agreement with existing experimental studies, the model shows that in attentional learning and multiple-cue category learning, both medicated and unmedicated PD patients are relatively impaired compared to healthy controls.

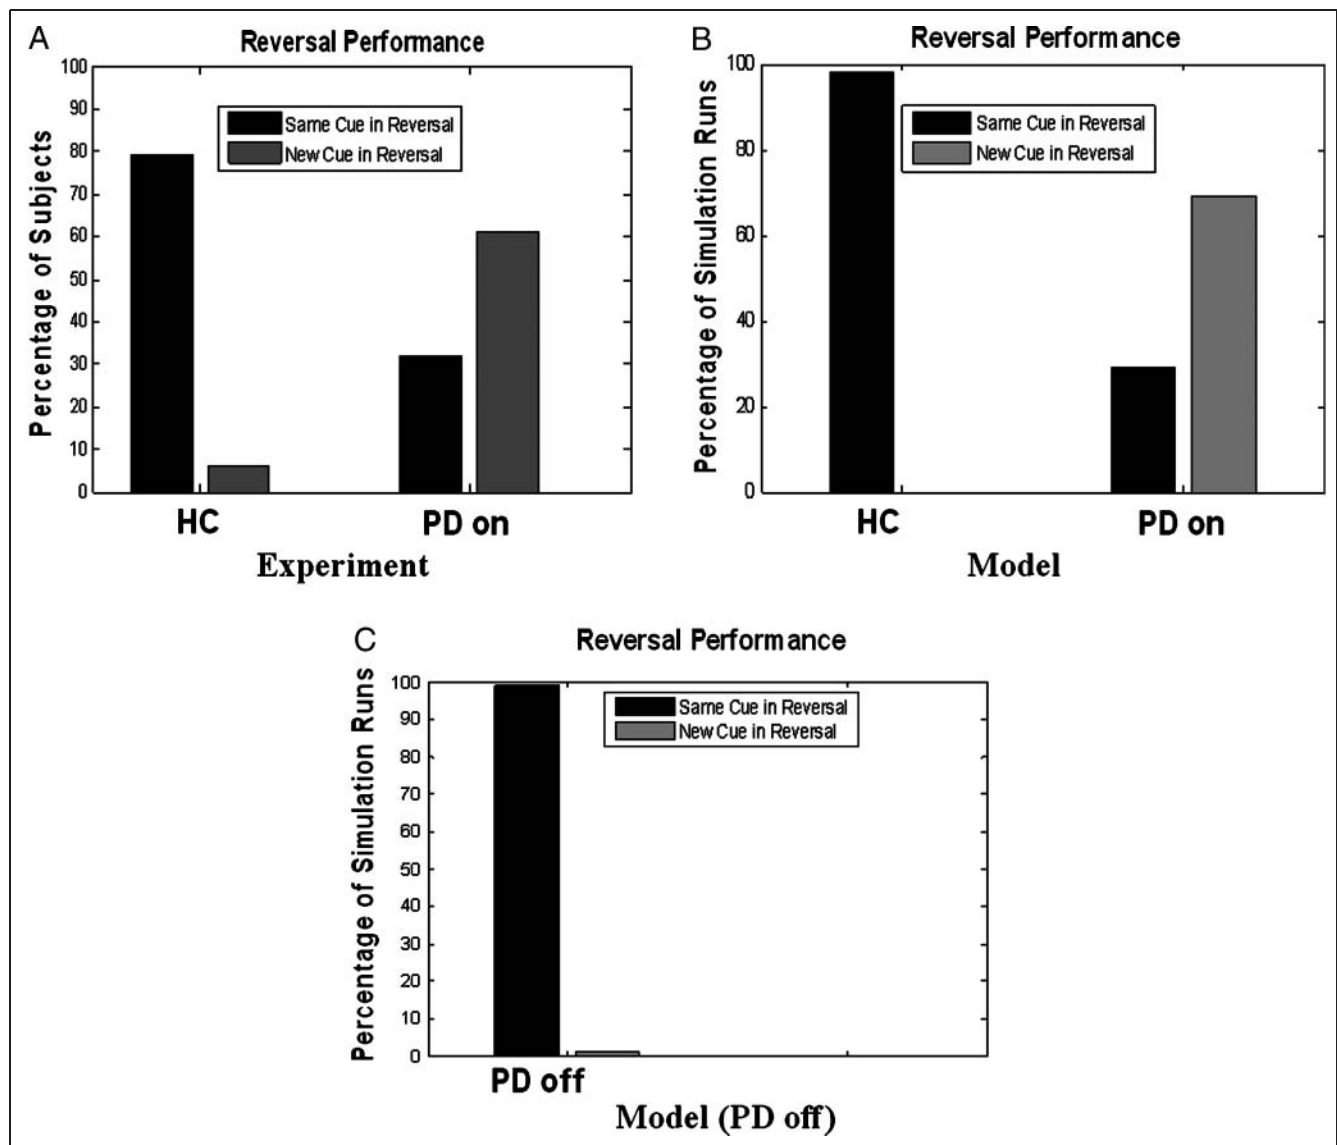
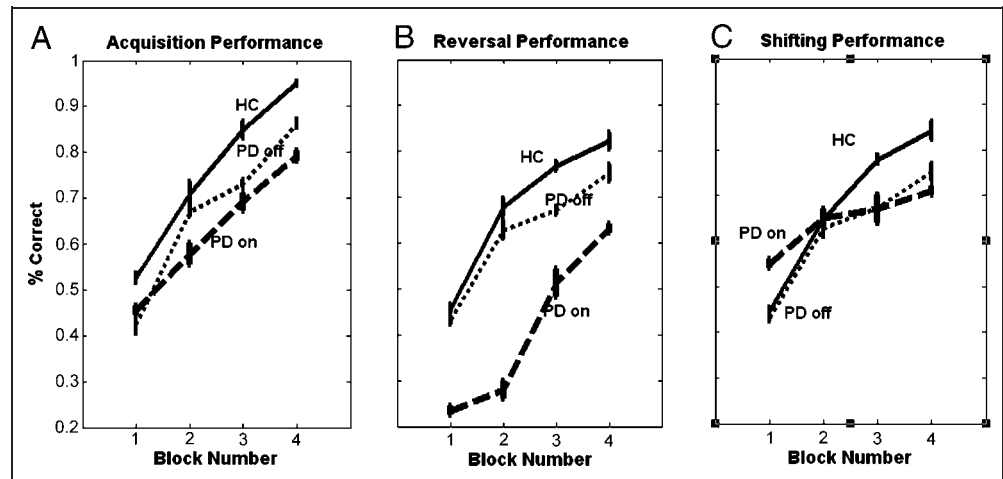


Figure 8. Performance in the reversal phase of the multicue category learning task (Shohamy et al., 2009). Here we assume that each run of the model represents a simulated subject performing the task, and thus, different runs of the model represent different subjects running the task. (A) Experimental results of the reversal phase of the multicue category learning task (adapted from Shohamy et al., 2009). Compare to simulation results (above). (B) Simulation results of the reversal phase of the multicue category learning task (Shohamy et al., 2009). Simulated HC: Like Shohamy et al. (2009), simulated HC subjects reversed along the same cue chosen during acquisition, although some few subjects opted out of reversal, which we did not simulate. Simulated PD on: Simulated medicated PD patients opted out of reversal in almost two-thirds of the runs (each run simulates a different PD patient). In some runs, the model did not opt out because it was still paying attention to the same cue chosen during the acquisition phase. This, however, happens in roughly 30% of the runs, which is in agreement with the experimental results (Shohamy et al., 2009). (C) Simulated unmedicated patients (PD off) reversed along the same cue chosen during acquisition, as in healthy controls.

Figure 9. Simulation results of the modified versions of the Shohamy et al. (2009) slot machine task in which only one cue is diagnostic of category membership (see Table 2 for task description). The acquisition phases in both forced-cue and shifting tasks are the same so they are presented once. (A) Acquisition performance in both forced-cue and shifting tasks. (B) Reversal performance of the forced-cue task. (C) Shifting performance in the shifting slot machine task. Unlike the original task, we found our simulated medicated PD patients to

be more impaired at reversal learning in the forced-cue task (see figure above) as compared to the original version of the task discussed previously. This is due in the model to increased attentional shifting performance in simulated medicated PD patients, which impairs performance on this task. In the shifting task, we found simulated medicated PD patients to be better than controls and unmedicated patients in the first block of the shifting phase. This is due to enhanced shifting in medicated PD patients. In the shifting task, we predict a significant interaction between block number and group, such that medicated patients are more enhanced in early block but more impaired toward the end of the shifting phase.



We first tested the model by simulating an instrumental conditioning task, which is considered a simple motor learning task. This is to assure us that adding a stimulus selection module to the model does not interfere with learning simple tasks. In addition, in agreement with many of the existing data, the model shows that both medicated and unmedicated PD patients should be relatively impaired at instrumental conditioning learning compared to controls, a finding supported by multiple experimental studies (Filoteo et al., 2007; Shohamy et al., 2006; Czernecki et al., 2002).

The model also simulated performance in the “weather prediction” task. In agreement with experimental studies, the model shows that (a) most subjects reach suboptimal (above chance) performance in this task and (b) medicated PD patients are more impaired than unmedicated PD patients and controls in this task, as found experimentally (Jahanshahi et al., 2010). Existing models of “weather prediction” task performance do not account for this effect (Frank, 2005). Finally, the model provides a qualitative fit to the multicue category task used by Shohamy et al. (2009). As in the “weather prediction” task, the model accounts for the finding that many subjects performing the “slot machine” task achieve only suboptimal performance. According to the model, suboptimal behavior in this task is related to attentional limitations, that is, paying attention only to one single cue of the presented pattern. One difference between the “weather prediction” task and the “slot machine” task is that, in the latter, each cue is equally diagnostic of the category membership, whereas in the “weather prediction” task, some cues are better predictors of category membership than others. Finally, the model provided a mechanistic account for why medicated PD patients opt out of reversal on the multicue category task of

Shohamy et al. An increase in tonic dopamine levels within PFC helps the model shift attention to other cues, which is in agreement with the verbal description by Cools et al. (2001) that dopaminergic medications enhance lateral PFC functioning but impair and overdose orbito-frontal cortex functioning.

Along the same lines, Cools et al. (2001) found that medicated PD patients show better performance at task switching than unmedicated PD patients. Aron, Poldrack, and Wise (2009) presented a qualitative model showing how PFC is key for task switching behavior. The findings that medicated PD patients (a) opted out of reversal during the performance of multiple-cue category learning tasks (Shohamy et al., 2009), (b) showed fewer perseverative errors than unmedicated PD patients in attentional shifting tasks (Rutledge et al., 2009; Owen et al., 1993), and (c) showed improved performance in the shifting version of the “slot machine” task (see Figure 9) can be accounted for in our model by a single mechanism. In all cases, medicated PD patients quickly learn to change rules and pay attention to different cues than one used during previous task phases. Another factor underlying the opt-out performance in medicated patients is that their performance in the acquisition phase was numerically lower than that of controls. In our model, this corresponds to lower attentional weight values than that of healthy controls (Figure 7). This makes it easier for medicated PD patients to switch attention to a different cue in the reversal phase. Accordingly, our model shows that overtraining medicated PD patients in the acquisition phase of Shohamy et al.’s (2009) task might lead to enhanced performance in the acquisition phase and fewer medicated patients will opt out of reversal than found in the original Shohamy et al. study, which is a new prediction of the model. Unlike the original

“slot machine” task, in the forced-cue task (Table 2), medicated PD patients were much more impaired at performing the reversal phase compared to unmedicated patients. This is due to enhanced attentional shifting, which impairs performance in this task. Simulation results of the original “slot machine” and forced-cue tasks might explain why different experimental studies report impairment (Czernecki et al., 2002; Cools et al., 2001) or enhancement (Rutledge et al., 2009) in reversal tasks in medicated PD patients. It might be the case that medicated PD patients can correctly perform reversal tasks that involve the option to shift attention to different cues (i.e., to perform the reversal phase using a different cue than ones used during previous acquisition phases).

PFC in the model is key for stimulus selection and attentional shifting. Disrupting PFC activity by adding noise leads to delayed learning and the occurrence of more perseverative errors than in the intact model. We simulated lesions by adding noise to the activation levels of PFC, as done in other models (Joanisse & Seidenberg, 1999; Olson & Humphreys, 1999). This is in agreement with behavioral observations (Owen et al., 1993) and modeling results (Amos, 2000).

Importantly, the model also assumes that PD is associated with decreased levels of both phasic and tonic dopamine levels in the basal ganglia and PFC. Dopamine medications, according to the model, increase tonic dopamine levels in the basal ganglia and PFC, and thus, lead to a smaller phasic dopamine signal than that of unmedicated PD patients. This is in agreement with existing experimental data showing that medicated patients are impaired at associative learning compared to unmedicated patients (Jahanshahi et al., 2010; Shohamy et al., 2006). This hypothesis is also in agreement with earlier models (Guthrie et al., 2009), although these models did not address the role of PFC dopamine to behavioral performance. This is, however, different from other computational models, which assumes that PD and dopaminergic medications mainly affect dopamine levels in the basal ganglia (Frank, 2005).

Existing Computational Models of Attentional and Motor Learning

We first review various neural network models that simulate performance in stimulus selection and attentional tasks. Second, we review basal ganglia models of action selection learning.

Kruschke (2003, 2005) developed a series of computational models that simulate the role of attention in the performance of different learning and attentional shifting paradigms. The main feature of Kruschke’s models is the inclusion of an attentional module that influences S–R associative learning. This module controls which perceptual information affects motor responses. In these models, activation of the attentional module is modulated by feedback provided to the network. By building different

models that revolve around this idea, Kruschke simulated performance in the highlighting paradigm, a modified blocking task that tests for attentional processes, and the intradimensional/extradimensional shifting (IDS/EDS) tasks (Kruschke, 2005). Similar to Kruschke’s models, our model assumes that corrective feedback affects stimulus learning, yet our model additionally maps brain processes.

Cohen and Servan-Schreiber (1992) provided a simple computational model of performance in the Stroop test. The main feature of this model is that active maintenance of the relevant dimension in PFC biases responses to perceptual information in favor of that dimension. Similarly, Cohen’s model assumes that amphetamine enhances activity and maintenance of relevant information in PFC. Dopamine projected to the attentional module was assumed to increase the signal-to-noise ratio. Recently, Aron et al. (2009) showed that a similar theory can account for performance in the task switching task, although it was not a simulation model. Furthermore, Stafford and Gurney (2007) have proposed a simulation model augmenting the Cohen model, in which they showed that adding a basal ganglia module to the original frontal module of Cohen’s models provides a better fit to human data on the Stroop test. These models do not incorporate the role of PFC dopamine to behavioral performance. However, similar to our model, Gurney and et al. have also argued that the basal ganglia is key for action selection and motor learning (Prescott et al., 2006).

Ashby et al. (1998) provided a conceptual model (termed COVIS) which assumes that the basal ganglia is key for implicit category learning, whereas PFC is key for verbal (explicit) category learning. Unlike Ashby et al.’s model, neither our model nor the experimental data we have simulated made a distinction between the performance of explicit and implicit category learning tasks. Similar to our model, Ashby et al. also assumed that PFC dopamine is key for rule selection, and that PD is associated with a decrease in dopamine levels in the basal ganglia and PFC.

Monchi, Taylor, and Dagher (2000) and Monchi and Taylor (1999) provided a biologically inspired computational model that simulates the role of the basal ganglia, thalamus, and PFC in different WM tasks. These models assume that basal ganglia input to PFC is key for maintenance of information in WM. Monchi et al. simulate PD by decreasing values of weights connecting PFC and striatal units. Unlike our model, Monchi et al. did not simulate the role of dopamine in learning, as suggested by experimental studies (Schultz et al., 1997). A potential extension of our model is to include basal ganglia–PFC loops which, according to Monchi et al., are key for WM performance.

Braver and Cohen (2000) provided a computational model showing how PFC learns to actively maintain relevant information in WM. The model was trained using the TD model, which accounts for experimental findings related to learning Pavlovian conditioning tasks. The design of the TD algorithm is based on the assumption that

reward prediction failure drives learning. The key feature of the Braver and Cohen model is that a positive TD error signal enhances the learning of gating of input into WM. Accordingly, information that is not associated with dopamine phasic signal, such as background noise or task-irrelevant stimuli, will not be gated into WM. This model simulates performance in the standard AX-CPT task, which was used in various experimental studies (Moustafa et al., 2008; Cohen, Barch, Carter, & Servan-Schreiber, 1999). In this task, subjects are presented with one stimulus at a time, some of which are relevant and some are not. Relevant stimuli are, by definition, associated with receiving positive feedback, and thus, elicit dopamine phasic responses. Accordingly, task-relevant stimuli are gated and maintained in WM in PFC (for more elaboration on these ideas, see Cohen, Braver, & Brown, 2002). Unlike our model, this model does not simulate stimulus selection processes.

Unlike the Braver and Cohen model, Hazy, Frank, & O'Reilly (2007) and O'Reilly and Frank (2006) argued that the basal ganglia is key for filtering out irrelevant information. Specifically, O'Reilly et al. argued that the basal ganglia is key for gating perceptual information into WM, whereas PFC is key for maintenance of information in WM. Based on the models of O'Reilly and Houk (1995), Moustafa and Maida (2007) show that the TD learning algorithm can be used to simulate S-R learning as well as learning to gate perceptual information into WM in delayed-response tasks.

Most models assume that PD and dopamine medications mainly affect basal ganglia processes (Guthrie et al., 2009; Moustafa & Maida, 2007; Frank, 2005), although experimental studies found evidence that dopamine medications do increase dopamine levels in PFC. Most importantly, the abovementioned models do not, however, simulate performance in attentional tasks, in which subjects should learn to pay attention to some stimuli and ignore co-occurring irrelevant stimuli or noise. This is a key problem for animal and human learning and survival as animals should be able to learn to identify relevant objects—or features of objects—that are important for survival. An interesting example given by Kruschke (2005) is that if flat-headed mushrooms are poisonous, but round-headed mushrooms are nutritious, one should learn to pay attention only to such relevant features of mushrooms (i.e., head shape) and ignore others (e.g., stalk) in order to be able to survive. The Braver and Cohen model does not simulate performance in such behavioral processes. A laboratory equivalent of such behavioral processes is the Wisconsin Card Sorting Test and some of the tasks used by Shepard, Hovland, and Jenkins (1961), in which paying attention to a subset of features of a presented pattern does sufficiently lead to optimal performance. Another similar paradigm is the “slot machine” task (Shohamy et al., 2009), in which paying attention to one cue of a presented multicue pattern leads to sufficiently good, although not optimal, behavior. The abovementioned models also did not incorporate the findings that dopamine medications increase dopamine levels in PFC, as found in experimental studies (Carey et al., 1995).

Limitations and Future Directions

The proposed model, although it can account for various multiple-cue category learning processes, has some limitations. For example, several studies have shown that dopaminergic medications enhance learning from positive feedback but impair learning from negative feedback (Frank, 2005). We do believe that an increase in tonic dopamine levels, as proposed here, can account for this reward learning effects. Specifically, an increase in tonic dopamine levels most likely increases excitability of spiny neurons in the striatum (Mallet, Ballion, Le Moine, & Gonon, 2006) and might thus facilitate learning in the direct cortico-striatal pathway through the three-factor learning rule (Reynolds et al., 2001). With regards to negative reinforcement or reversal deficits in medicated PD patients, it is perhaps the case that dopamine medications (a) overdose orbito-frontal cortex or inferior frontal cortex, which are brain regions that were shown to be key for inhibiting motor responses (Cools et al., 2001), and/or (b) enhance learning in the basal ganglia indirect pathway (Frank, 2005). Future models will attempt to address these studies, and address why some studies report reversal deficits in medicated PD patients (Cools et al., 2001), whereas others do not report reversal deficits (Rutledge et al., 2009).

Furthermore, the model, in its current form, can only learn to pay attention to one cue at a time. This, we recognize, is a highly oversimplified model of stimulus selection, as animals and human subjects can learn complex tasks that require paying attention to configurations of multiple cues, such as in the negative patterning task (Bussey et al., 2000), or tasks used by Shepard et al. (1961). Learning these complex tasks has been found to engage additional brain areas, including the anterior cingulate. Furthermore, some argue that the hippocampal region is needed for configural learning, and thus, the use of multicue category learning (Rudy & Sutherland, 1989). According to the Gluck and Myers (1993) model, the hippocampus learns to form compressed and differentiated representations of input stimuli. Accordingly, in order for the model to reach optimal behavior and employ a multicue strategy in both the “weather prediction” task and Shohamy et al.’s (2009) task, it will learn to form new representations of all presented patterns and then learn to map these new representations to various responses. Like most actor-critic models (Suri & Schultz, 1999; Houk, 1995), the model presented here only learns to make one motor response at a time. Along the same lines, our model does not account for transfer generalization processes in acquisition and acquired equivalence tasks because these processes arguably rely on the integrity of the hippocampal region (Myers et al., 2002; Gluck & Myers, 1993). Future modeling plans will attempt to simulate the role of the anterior cingulate and hippocampus in more demanding attentional processes.

It is important to note that a good fit does not make a theory plausible (for a discussion, see Roberts & Pashler, 2000): A good fit is but one aspect of a plausible theory.

The possibility that some other parameter values might work (which might then suggest a different theory) is unavoidable (see Roberts & Pashler, 2000). We also found that, assuming that PD and dopamine medications affect (a) tonic dopamine levels alone or (b) affect one brain structure alone, did not account for all existing data, although our attempts do not rule out the possibility of finding other parameter values (and theories) to explain existing results.

Furthermore, in addition to providing a good fit, our model is largely constrained by existing biological data (e.g., role of dopamine in learning and performance, effects of dopamine medication on different brain areas, among others). One limitation, however, is that like most existing models (Frank, 2005; Amos, 2000; Suri & Schultz, 1999; Gluck & Myers, 1993), values chosen for learning rate and gain parameters in our model do not necessarily correspond to biological features, and it is possible that some other values might lead to the same results (see Roberts & Pashler, 2000 for a discussion). However, we found that actual values are not critical for simulation results presented above, that is, variations in parameter values do not dramatically change model performance. Nevertheless, the key feature of the model is that parameter values (or minor variations of these values) chosen for the simulations of PD and dopaminergic medications might reflect how PD and dopaminergic medications affect biological features (phasic and tonic dopamine levels). The model suggests that PD and dopaminergic medications affect phasic and tonic dopamine firing in the basal ganglia and PFC. Furthermore, we also note that a model's learning performance does not exactly match the learning performance of the experimental data. This is because the goal of the model is to capture behavioral differences (a) in various tasks and (b) among different groups (controls and two PD patient groups), which the model was able to simulate with some success.

Another limitation is that the model does not incorporate functional contributions of different dopamine receptors to behavioral performance. For example, Frank (2005) argued that D1 receptors in the basal ganglia are key for reward learning, whereas D2 receptors are key for punishment and avoidance learning (for similar ideas, see Prescott et al., 2006). As mentioned above, the model, in its current form, does not simulate performance in reward/punishment learning tasks. The Frank model does not address the differential contribution of different DA receptors in PFC. Cohen et al. (2002) argued that D1 and D2 receptors in PFC play different roles, such that tonic dopamine is key for maintenance of information in WM via D1 receptors modulation, whereas phasic dopamine is key for learning (synaptic modification) via D2 receptor modulation (but see Schultz, 2007 for different ideas regarding behavioral functions of D1 and D2 receptors). Our model does not incorporate the differential contributions of dopamine receptors to behavioral performance, although it does not necessarily conflict with assumptions of the Frank and Cohen models. It is important to note here that some studies found that increase of dopamine in PFC does im-

pair its function (Takahashi et al., 2008). We are led to assume that it is perhaps an excessive increase in dopamine levels beyond that of normal doses of Parkinsonian medications that might impair prefrontal function. This finding is not accounted for by our model.

Another limitation of the model is that it only provides a qualitative fit to existing behavioral data. This is because the model has a few free parameters and, within this constraint, seeks to account for a broad range of experimental data. To provide a quantitative fit to these and other behavioral results, a model would require many more free parameters.

Despite its limitations, our model does provide several new predictions. For example, in the multiple-cue category learning task used by Shohamy et al. (2009), we predict that nonmedicated PD patients perform better than medicated PD patients at learning the acquisition phase. We also predict that unmedicated PD patients, unlike medicated PD patients, will reverse successfully along the same cue chosen during the acquisition phase. As mentioned above, we also predict that overtraining medicated PD patients on the Shohamy et al. (2009) "slot machine" task might lead to enhanced performance in the acquisition phase, and that fewer medicated PD patients will then opt out of reversal than in the original Shohamy et al. study. Furthermore, as shown in Figure 9, we predict that medicated PD patients will show impairment at the reversal phase of the forced-cue "slot machine" task but will show enhancement at the shifting phase of the modified "slot machine" task (see Table 2 for task description). Future experimental work should confirm or refute these predictions.

APPENDIX

Model Details: Learning and Weight Update

We assume that learning in the attentional (PFC) and motor (striatal) modules relies on phasic dopamine signals projected from the midbrain (for similar ideas, see Suri & Schultz, 1999). In this model, phasic dopamine signals are key for both attentional and S-R learning. The model is trained using the TD algorithm, which simulates various characteristics of phasic dopamine firing (Schultz et al., 1997). Let $TD(t)$ be the TD error signal at time t (also known as the effective reinforcement); $R(t)$ be the reward presented at time t (reward is 1 when reward is presented after correct feedback and is 0 otherwise); $P(t)$ be the reward prediction at time t ; γ be the discount factor (which determines how future reward affect reward predictions; is set to 0.99 in all simulation runs presented here). The TD error is computed as follows:

$$TD(t) = R(t) + \gamma P(t) - P(t - 1)$$

Let w_i be the weight connecting unit i to the critic node, n be the number of input nodes, and x_i be activation of

input units (which take binary (0,1) values). Reward prediction $P(t)$ is computed by the critic node as follows:

$$P(t) = \sum_{i=1}^n w_i(t)x_i(t)$$

Now, we describe the equations of the actor module. Let w_{ji} be the weight connecting unit i to unit j ; $\delta_{ji}(t)$ be the Gaussian noise associated with the weight w_{ji} (with zero mean and standard deviation of 0.025; also see Moustafa & Maida, 2007).

All weights are perturbed using Gaussian noise, which is included to induce exploration in the system. Let u_{ji} be the perturbed weight connecting unit i to unit j . Perturbed weight values are computed as follows:

$$u_{ji}(t) = w_{ji}(t) + \delta_{ji}(t),$$

Activations of all units in the network are computed using a sigmoidal function:

$$f(x) = \frac{1}{1 + e^{-G_m x}}$$

where G_m is the gain parameter (G_{pfc} for PFC and G_{bg} for basal ganglia module). Let n be the number of input (or prefrontal) units. The input units take binary (0,1) values. The activation of a unit j is computed as:

$$A_j(t) = f\left(\sum_{i=1}^n u_{ji}(t)x_i(t)\right)$$

In the model, a winner-take-all network computes the unit with the highest activation in both the attentional (PFC) and motor (striatal) modules. In other words, we assume that a winner-take-all competition among striatal neurons is the mechanism underlying the choice of motor responses. Similarly, we assume that a winner-take-all competition among PFC neurons is the mechanism underlying competition among *attentional* nodes.

$$A_j^p = \begin{cases} 1 & \text{if } A_j > \beta \text{ \& } A_j > A_i \text{ for all } i \neq j \\ 0 & \text{otherwise} \end{cases}$$

where β is a threshold, A_j is the activation of unit j , A_j^p is the activation of unit j resulting from winner-take-all computations (for similar ideas, see Suri & Schultz, 1999; Schultz et al., 1997).

Learning in the model is based on the three-factor rule of learning—also known as the dopamine-based Hebbian learning rule (for similar ideas, see Guthrie et al., 2009). According to this rule, the phasic dopamine signal is key for strengthening weights linking active nodes. It is also key for

weakening weights linking an active node and another inactive node. Also, different computational models incorporate this learning rule (Guthrie et al., 2009; Braver & Cohen, 2000; Suri & Schultz, 1999).

Let LR_m be the learning rate. There are two learning rate parameters in the model, one for the PFC (stimulus selection) module and one for the basal ganglia (motor) response module. Let x_i represent the activation level of the presynaptic node. The weight update rule is

$$w_{ji}(t+1) = w_{ji}(t) + \text{LR}_m \text{TD}(t)x_i(t)A_j^p$$

Lesioning PFC in the model is simulated by adding noise to activation levels of PFC nodes. Let δ be Gaussian noise, A_j^{pre} be activation before adding noise to PFC activation levels, A_j^{post} be activation after adding noise to PFC activation. Then, for every PFC unit j ,

$$A_j^{\text{post}} = A_j^{\text{pre}} + \delta_j$$

Strategy Analysis

We have conducted strategy analysis on model output from the “slot machine” task following the same procedures described previously for use with empirical data from human subjects (Shohamy et al., 2004, 2009; Gluck et al., 2002). In the original “slot machine” task, suboptimal performance (80% correct) can be achieved using a one-cue strategy based on any one of the three cues, whereas optimal performance (100% correct) can be achieved using a configural strategy based on all three cues. To determine which strategy a subject was using, that subject’s trial-by-trial responses were compared against “ideal data” that would have been expected if a subject were consistently following either a one-cue or a configural strategy. In the same way, we compared the model output in each simulation run against ideal data and classed the simulation’s “strategy,” based on which set of ideal data more closely approximated the model’s actual performance. Most importantly, in the reversal phase, we considered whether, in each run, the model was responding based on the same specific cue during acquisition and reversal, but simply reversed response valence, or whether responding was based on a new cue during the reversal, meaning attention had shifted. This approach allowed us to classify individual simulation runs as attending to the same cue during acquisition and reversal, or as shifting attention to a new cue during the reversal phase.

Acknowledgments

We thank Catherine Myers, Arda Hotz, Daphna Shohamy, Martin Guthrie, and Martijn Meeter for helpful and insightful comments on the model details and earlier drafts of the manuscript.

Funding

Dekker Foundation; Bachman-Strauss Dystonia and Parkinson Foundation; NIH-NINDS (R01-NS047434-02 to M. A. G.). This work was partially supported by a grant to M. A. G. from the National Science Foundation (BCS-022 3910).

Reprint requests should be sent to Ahmed A. Moustafa, Center for Molecular and Behavioral Neuroscience, Rutgers University, Room 209, 197 University Ave., Newark, NJ 07102, or via e-mail: ahmedhalimo@gmail.com.

REFERENCES

- Amos, A. (2000). A computational model of information processing in the frontal cortex and basal ganglia. *Journal of Cognitive Neuroscience*, *12*, 505–519.
- Aron, A. R., Poldrack, R. A., & Wise, S. P. (2009). Cognition: Basal ganglia role. In L. R. Squire (Ed.), *Encyclopedia of neuroscience* (Vol. 2, pp. 1069–1077). Oxford: Academic Press.
- Ashby, F. G., Alfonso-Reese, L. A., Turken, A. U., & Waldron, E. M. (1998). A neuropsychological theory of multiple systems in category learning. *Psychological Review*, *105*, 442–481.
- Ashby, F. G., Ell, S. W., Valentin, V. V., & Casale, M. B. (2005). FROST: A distributed neurocomputational model of working memory maintenance. *Journal of Cognitive Neuroscience*, *17*, 1728–1743.
- Braver, T. S., & Cohen, J. D. (2000). On the control of control: The role of dopamine in regulating prefrontal function and working memory. In S. Monsell & J. Driver (Eds.), *Control of cognitive processes: Attention and performance XVIII* (chapter 16, 779 pp.). Cambridge, MA: MIT Press.
- Bussey, T. J., Dias, R., Redhead, E. S., Pearce, J. M., Muir, J. L., & Aggleton, J. P. (2000). Intact negative patterning in rats with fornix or combined perirhinal and postrhinal cortex lesions. *Experimental Brain Research*, *134*, 506–519.
- Carey, R. J., Pinheiro-Carrera, M., Dai, H., Tomaz, C., & Huston, J. P. (1995). L-DOPA and psychosis: Evidence for L-DOPA-induced increases in prefrontal cortex dopamine and in serum corticosterone. *Biological Psychiatry*, *38*, 669–676.
- Cohen, J. D., Barch, D. M., Carter, C., & Servan-Schreiber, D. (1999). Context-processing deficits in schizophrenia: Converging evidence from three theoretically motivated cognitive tasks. *Journal of Abnormal Psychology*, *108*, 120–133.
- Cohen, J. D., Braver, T. S., & Brown, J. W. (2002). Computational perspectives on dopamine function in prefrontal cortex. *Current Opinion in Neurobiology*, *12*, 223–229.
- Cohen, J. D., & Servan-Schreiber, D. (1992). Context, cortex, and dopamine: A connectionist approach to behavior and biology in schizophrenia. *Psychological Review*, *99*, 45–77.
- Cools, R., Barker, R. A., Sahakian, B. J., & Robbins, T. W. (2001). Enhanced or impaired cognitive function in Parkinson's disease as a function of dopaminergic medication and task demands. *Cerebral Cortex*, *11*, 1136–1143.
- Czernecki, V., Pillon, B., Houeto, J. L., Pochon, J. B., Levy, R., & Dubois, B. (2002). Motivation, reward, and Parkinson's disease: Influence of dopatherapy. *Neuropsychologia*, *40*, 2257–2267.
- Daw, N. D., Niv, Y., & Dayan, P. (2005). Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nature Neuroscience*, *8*, 1704–1711.
- Fera, F., Weickert, T. W., Goldberg, T. E., Tessitore, A., Hariri, A., Das, S., et al. (2005). Neural mechanisms underlying probabilistic category learning in normal aging. *Journal of Neuroscience*, *25*, 11340–11348.
- Filoteo, J. V., Maddox, W. T., Ing, A. D., & Song, D. D. (2007). Characterizing rule-based category learning deficits in patients with Parkinson's disease. *Neuropsychologia*, *45*, 305–320.
- Frank, M. J. (2005). Dynamic dopamine modulation in the basal ganglia: A neurocomputational account of cognitive deficits in medicated and nonmedicated Parkinsonism. *Journal of Cognitive Neuroscience*, *17*, 51–72.
- 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, U.S.A.*, *104*, 16311–16316.
- Gluck, M. A., & Myers, C. E. (1993). Hippocampal mediation of stimulus representation: A computational theory. *Hippocampus*, *3*, 491–516.
- Gluck, M. A., Shohamy, D., & Myers, C. (2002). How do people solve the “weather prediction” task?: Individual variability in strategies for probabilistic category learning. *Learning and Memory*, *9*, 408–418.
- Guthrie, M., Myers, C. E., & Gluck, M. A. (2009). A neurocomputational model of tonic and phasic dopamine in action selection: A comparison with cognitive deficits in Parkinson's disease. *Behavioural Brain Research*, *200*, 48–59.
- Hazy, T. E., Frank, M. J., & O'Reilly, R. C. (2007). Towards an executive without a homunculus: Computational models of the prefrontal cortex/basal ganglia system. *Philosophical Transactions of the Royal Society of London, Series B, Biological Sciences*, *362*, 1601–1613.
- Houk, J. C. (1995). Information processing in modular circuits linking basal ganglia and cerebral cortex. In J. C. Houk, J. L. Davis, & D. G. Beiser (Eds.), *Models of information processing in the basal ganglia* (chapter 12, 382 pp.). Cambridge, MA: MIT Press.
- Iba, M., & Sawaguchi, T. (2003). Involvement of the dorsolateral prefrontal cortex of monkeys in visuospatial target selection. *Journal of Neurophysiology*, *89*, 587–599.
- Jahanshahi, M., Wilkinson, L., Gahir, H., Dharminda, A., & Lagnado, D. A. (2010). Medication impairs probabilistic classification learning in Parkinson's disease. *Neuropsychologia*, *48*, 1096–1103.
- Jellinger, K. A. (1999). Post mortem studies in Parkinson's disease: Is it possible to detect brain areas for specific symptoms? *Journal of Neural Transmission. Supplementum*, *56*, 1–29.
- Joanisse, M. F., & Seidenberg, M. S. (1999). Impairments in verb morphology after brain injury: A connectionist model. *Proceedings of the National Academy of Sciences, U.S.A.*, *96*, 7592–7597.
- Kaasinen, V., Nurmi, E., Bergman, J., Eskola, O., Solin, O., Sonninen, P., et al. (2001). Personality traits and brain dopaminergic function in Parkinson's disease. *Proceedings of the National Academy of Sciences, U.S.A.*, *98*, 13272–13277.
- Kish, S. J., Shannak, K., & Hornykiewicz, O. (1988). Uneven pattern of dopamine loss in the striatum of patients with idiopathic Parkinson's disease. Pathophysiologic and clinical implications. *New England Journal of Medicine*, *318*, 876–880.
- Knowlton, B. J., Mangels, J. A., & Squire, L. R. (1996). A neostriatal habit learning system in humans. *Science*, *273*, 1399–1402.
- Kruschke, J. K. (2003). Attention in learning. *Current Directions in Psychological Science*, *12*, 171–175.
- Kruschke, J. K. (2005). Learning involves attention. In G. Houghton (Ed.), *Connectionist models in cognitive psychology* (pp. 113–140). East Sussex, UK: Psychology Press.

- Lewis, S. J., Slabosz, A., Robbins, T. W., Barker, R. A., & Owen, A. M. (2005). Dopaminergic basis for deficits in working memory but not attentional set-shifting in Parkinson's disease. *Neuropsychologia*, *43*, 823–832.
- Mallet, N., Ballion, B., Le Moine, C., & Gonon, F. (2006). Cortical inputs and GABA interneurons imbalance projection neurons in the striatum of parkinsonian rats. *Journal of Neuroscience*, *26*, 3875–3884.
- Monchi, O., & Taylor, J. G. (1999). A hard wired model of coupled frontal working memories for various tasks. *Information Sciences*, *113*, 221–243.
- Monchi, O., Taylor, J. G., & Dagher, A. (2000). A neural model of working memory processes in normal subjects, Parkinson's disease and schizophrenia for fMRI design and predictions. *Neural Networks*, *13*, 953–973.
- Moustafa, A. A., & Maida, A. S. (2007). Using TD learning to simulate working memory performance in a model of the prefrontal cortex and basal ganglia. *Cognitive Systems Research*, *8*, 262–281.
- Moustafa, A. A., Sherman, S. J., & Frank, M. J. (2008). A dopaminergic basis for working memory, learning and attentional shifting in Parkinsonism. *Neuropsychologia*, *46*, 3144–3156.
- Myers, C. E., Kluger, A., Golomb, J., Ferris, S., de Leon, M. J., Schnirman, G., et al. (2002). Hippocampal atrophy disrupts transfer generalization in nondemented elderly. *Journal of Geriatric Psychiatry and Neurology*, *15*, 82–90.
- Olson, A. C., & Humphreys, G. W. (1999). Connectionist models of neuropsychological disorders. *Trends in Cognitive Sciences*, *1*, 222–228.
- O'Reilly, R. C., & Frank, M. J. (2006). Making working memory work: A computational model of learning in the frontal cortex and basal ganglia. *Neural Computation*, *18*, 283–328.
- Owen, A. M., Roberts, A. C., Hodges, J. R., Summers, B. A., Polkey, C. E., & Robbins, T. W. (1993). Contrasting mechanisms of impaired attentional set-shifting in patients with frontal lobe damage or Parkinson's disease. *Brain*, *116*, 1159–1175.
- Prediger, R. D., Batista, L. C., Medeiros, R., Pandolfo, P., Florio, J. C., & Takahashi, R. N. (2006). The risk is in the air: Intranasal administration of MPTP to rats reproducing clinical features of Parkinson's disease. *Experimental Neurology*, *202*, 391–403.
- Prescott, T. J., Montes Gonzalez, F. M., Gurney, K., Humphries, M. D., & Redgrave, P. (2006). A robot model of the basal ganglia: Behavior and intrinsic processing. *Neural Networks*, *19*, 31–61.
- Reynolds, J. N., Hyland, B. I., & Wickens, J. R. (2001). A cellular mechanism of reward-related learning. *Nature*, *413*, 67–70.
- Roberts, S., & Pashler, H. (2000). How persuasive is a good fit? A comment on theory testing. *Psychological Review*, *107*, 358–367.
- Rudy, J. W., & Sutherland, R. J. (1989). The hippocampal formation is necessary for rats to learn and remember configural discriminations. *Behavioural Brain Research*, *34*, 97–109.
- Rutledge, R. B., Lazzaro, S. C., Lau, B., Myers, C. E., Gluck, M. A., & Glimcher, P. W. (2009). Dopaminergic drugs modulate learning rates and perseveration in Parkinson's patients in a dynamic foraging task. *Journal of Neuroscience*, *29*, 15104–15114.
- Schultz, W. (2007). Multiple dopamine functions at different time courses. *Annual Review of Neuroscience*, *30*, 259–288.
- Schultz, W., Dayan, P., & Montague, P. R. (1997). A neural substrate of prediction and reward. *Science*, *275*, 1593–1599.
- Servan-Schreiber, D., Bruno, R. M., Carter, C. S., & Cohen, J. D. (1998). Dopamine and the mechanisms of cognition: Part I. A neural network model predicting dopamine effects on selective attention. *Biological Psychiatry*, *43*, 713–722.
- Servan-Schreiber, D., Cohen, J. D., & Steingard, S. (1996). Schizophrenic deficits in the processing of context. A test of a theoretical model. *Archives of General Psychiatry*, *53*, 1105–1112.
- Shepard, R. N., Hovland, C. I., & Jenkins, H. M. (1961). Learning and memorization of classifications. *Psychological Monographs*, *75*.
- Shohamy, D., Myers, C. E., Kalanithi, J., & Gluck, M. A. (2008). Basal ganglia and dopamine contributions to probabilistic category learning. *Neuroscience and Biobehavioral Reviews*, *32*, 219–236.
- Shohamy, D., Myers, C. E., Gekhman, K. D., Sage, J., & Gluck, M. A. (2006). 1-Dopa impairs learning, but spares generalization, in Parkinson's disease. *Neuropsychologia*, *44*, 774–784.
- Shohamy, D., Myers, C. E., Hopkins, R. O., Sage, J., & Gluck, M. A. (2009). Distinct hippocampal and basal ganglia contributions to probabilistic learning and reversal. *Journal of Cognitive Neuroscience*, *21*, 1821–1833.
- Shohamy, D., Myers, C. E., Onlaor, S., & Gluck, M. A. (2004). Role of the basal ganglia in category learning: How do patients with Parkinson's disease learn? *Behavioral Neuroscience*, *118*, 676–686.
- Silberstein, P., Pogosyan, A., Kuhn, A. A., Hotton, G., Tisch, S., Kupsch, A., et al. (2005). Cortico-cortical coupling in Parkinson's disease and its modulation by therapy. *Brain*, *128*, 1277–1291.
- Stafford, T., & Gurney, K. N. (2007). Biologically constrained action selection improves cognitive control in a model of the Stroop task. *Philosophical Transactions of the Royal Society of London, Series B, Biological Sciences*, *362*, 1671–1684.
- Suri, R. E., & Schultz, W. (1999). A neural network model with dopamine-like reinforcement signal that learns a spatial delayed response task. *Neuroscience*, *91*, 871–890.
- Tadaiesky, M. T., Dombrowski, P. A., Figueiredo, C. P., Cargnin-Ferreira, E., Da Cunha, C., & Takahashi, R. N. (2008). Emotional, cognitive and neurochemical alterations in a premotor stage model of Parkinson's disease. *Neuroscience*, *156*, 830–840.
- Takahashi, H., Kato, M., Takano, H., Arakawa, R., Okumura, M., Otsuka, T., et al. (2008). Differential contributions of prefrontal and hippocampal dopamine D(1) and D(2) receptors in human cognitive functions. *Journal of Neuroscience*, *28*, 12032–12038.