Supporting Online Material for

Hold Your Horses: Impulsivity, Deep Brain Stimulation, and Medication in Parkinsonism

Michael J. Frank,* Johan Samanta, Ahmed A. Moustafa, Scott J. Sherman

*To whom correspondence should be addressed. E-mail: mfrank@u.arizona.edu

Published 25 October 2007 on *Science* Express
DOI: 10.1126/science.1146157

This PDF file includes:

Supporting Online Material
Hold Your Horses: Impulsivity, Deep Brain Stimulation and Medication in Parkinsonism

Michael J. Frank[1*], Johan Samanta [2,3], Ahmed Moustafa[1] & Scott J. Sherman[3]

*Corresponding author: mfrank@u.arizona.edu

[1] Dept of Psychology and Program in Neuroscience, University of Arizona

[2] Banner Good Samaritan Medical Center, Phoenix AZ

[3] Dept of Neurology, University of Arizona

1 Materials and Methods

Procedures were approved by the University of Arizona Human Subjects Protection Program.

1.1 Participants

We tested 17 DBS patients, 15 medication patients, and 27 healthy senior controls matched for age, education, and scores on the North American Adult Reading Test (NAART), an estimate of premorbid verbal IQ (*1*). The vast majority of patients performed the experiment twice, on and off medication or on and off DBS (order counterbalanced). DBS patients were medicated with their regular dopamine medication (at lower dosages) in both sessions, so that any effect can be attributed to DBS itself, and so that potential findings could be considered clinically relevant. A minority of patients opted to not return for a second session (1 patient on medication, 4 on DBS and 1 off DBS).

Because we were interested in the extent to which participants experienced conflict and/or learned about the positive versus negative outcomes of their choices, we had to first ensure that they learned the basic task. While the training criteria were meant to address this issue, some participants were globally confused by the lack of feedback and addition of novel pairs during test and therefore performed poorly all around, including in pairs which were easiest for them during training. As in previous studies, we eliminated participants from the analysis who did not perform better than chance (50%) during test in the easiest AB pair (*2*). This amounted to 7 single sessions from senior controls (2 of these seniors met criteria in the other session), 2 patients on medication, 5 patients off medication, 2 patients on DBS and 2 off DBS. The remaining participants included in analysis amounted to 22 senior controls, 12 patients off STN-DBS, 15 on STN-DBS, 12 on medication and 9 off medication.

The demographics of seniors and PD patients are shown in Table 1. All patients were receiving daily L-Dopa preparations, and virtually all of them were supplemented with D2 receptor agonists

(pramipexole or ropinrole) and/or selegiline (monoamine activity enhancer), and were stable on their medication dose for at least 2 months. Patients in the off medication condition withheld taking their regular dose of all DA-related medications for a mean of 19 hours prior to the experiment (minimum 15 hours) (*2*). Participating senior controls were either the spouses of PD patients or were recruited from local Tucson senior centers. Exclusionary criteria were as follows:

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

1.2 Deep Brain Stimulation Surgical Procedures

Patients underwent standard bilateral simultaneous deep brain stimulation surgery after clinical determination of the presence of levodopa-responsive Parkinson's disease which was no longer adequately treated by standard dopamine-replacement therapy. None of the patients tested had undergone prior ablative surgical procedures. A stereotactic frame (Leksell stereotactic system) was placed on the morning of the surgery and patient underwent imaging for direct targeting. The initial surgical target was determined by a combination of direct and indirect targeting using a neuronavigational computer workstation (Framelink, Medtronic). Indirect targeting utilized the default settings of the system (12mm lateral, 3mm posterior, 3 mm inferior to the middle point of the anterior commissural-posterior commissural line). These were modified if necessary by direct targeting primarily utilizing inversion recovery MRI. In some cases co-registration of CT and MRI were utilized, which also allows for a more accurate assessment of placement post operatively (via CT to composite comparison). Once the target was determined microelectrode recording under local anesthesia was utilized to refine the targeting. A tract was considered acceptable if the span of the STN transected was more than 5 mm and contained sensorimotor neuronal units. The inferior border of the STN was determined by locating the relatively silent region immediately superior to the substantia nigra pars reticulata identified by characteristic high frequency firing pattern. The chronic stimulation electrode (Medtronic 3387 or 3389) was then placed with the distal contact at the inferior border of the STN. The suitability of the placement was confirmed by intra-operative macrostimulation to ensure that there was improvement in the motor symptoms and absence of stimulation side effects at 4.0 volts, 60 microsecond pulse width and 185 Hz frequency with bipolar stimulation. Two weeks after electrode implantation, a neurogenerator was placed (Kinetra or Soletra, Medtronic).

Electrode placement was confirmed to be within the subthalamic region via post-operative MRI (see Figure 1 for a prototypical example).

1.3 Behavioral Methods

Participants were tested in two separate experimental sessions, separated by a minimum of 7 days, using different stimuli across sessions (*3*). Patients in the off medication session abstained from taking their regular dose of dopaminergic medications for a mean of 19 hours prior to the experiment (*2, 4*). For off DBS sessions, DBS units were turned off by the physician 30 minutes prior to the cognitive experiment. The order of off/on medication and off/on DBS was counterbalanced across patients. Following the end of the standard task, patients who had been off DBS then had their stimulators turned back on. A 10 minute delay ensued, after which a second test phase was administered. The same procedure was employed in a subset (N=14) of senior controls, without the DBS manipulation. Patients off medication were also tested in a second test phase, after taking a regular dose of their levodopa medication, but with a delay of 40-60 minutes (to ensure medication would have time to absorb).

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

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

We enforced a performance criterion (evaluated after each training block of 60 trials) to ensure that all participants were at the same performance level before advancing to test. Because of the different probabilistic structure of each stimulus pair, we used a different criterion for each (65% A in AB, 60% C in CD, 50% E in EF). (In the EF pair, stimulus E is correct 60% of the time, but this is particularly difficult to learn. We therefore used a 50% criterion for this pair simply to ensure that if participants happened to "like" stimulus F at the outset, they nevertheless had to learn that this bias was not going to consistently work.). The participant advanced to the test session if all these criteria were met, or after six blocks (360 trials) of training.

Participants were subsequently tested with the same training pairs, in addition to all novel combinations of stimuli, in random sequence. Prior to the test phase they were given the following instructions: "*It's time to test what you've learned! During this set of trials you will NOT receive feedback ('Correct' or 'Incorrect' to your responses. If you see new combinations of symbols in the test, please choose the symbol that 'feels' more correct based on what you learnt during the training sessions. If you're not sure which one to pick, just go with your gut instinct!*" Each

test pair was presented four times for a maximum of four seconds duration, and no feedback was provided.

When comparing high to low conflict RT's, we first collapsed high conflict win/win and lose/lose decisions and compared them to all low conflict pairs. Subsequent analyses compared high conflict win/win pairs to low conflict pairs involving a positive stimulus (AD and AF), whereas high conflict lose/lose pairs were compared with low conflict pairs involving a negative stimulus (BC and BE).

2 Additional Results and Analysis

2.1 Data Analysis

We used SAS v9.1 PROC MIXED to examine both between and within subject differences, using unstructured covariance matrices (which does not make any strong assumptions about the variance and correlation of the data, as do structured covariances). Where indicated, we tested for specific planned contrasts. In these contrasts, the number of degrees of freedom reflects the entire sample, and not just the participants involved in the particular contrast, because the mixed procedure analyzes both between and within effects. This procedure uses all of the data to provide a more stable estimate of the error term (*5*).

2.2 Accuracy during the Learning Phase

In the learning phase of the task, there was a trend for a main effect of group that neared significance ($F[4,51] = 2.5$, $p = .055$). Planned comparisons revealed that this was due to PD patients overall performing worse than seniors ($F[1,51] = 7.84$, $p = 0.0072$). This result is consistent with previous observations that PD patients have difficulty learning in probabilistic reinforcement paradigms, possibly due to depleted striatal dopamine (*6–8*). There was no difference in learning phase accuracy between patients on vs off DBS or on vs off medication, or in the DBS group vs the medication patients group (p's > 0.4). There were also no differences between these groups in the number of training trials to reach criterion before advancing to test phase (p's > 0.3), except that again PD patients as a whole were slower to do so than senior controls ($F[1,51] = 5.84$, $p = 0.02$). The mean (standard error) trials taken for each group was as follows: Seniors, 191 (25); off medication 267 (40), on medication 230 (32); off DBS 285 (31); on DBS 291 (28).

2.3 Accuracy in Positive versus Negative Test Choices

We replicated previously observed findings (*2*) in a different geographic setting, and in a within-subject design (all patients performed the same task on and off medication, whereas in the previous study different tasks were used between sessions). There was a significant interaction between medication status and choose-A versus avoid-B test performance ($F[1,51] = 4.1$, $p = .048$). Medicated patients performed significantly better at choose-A than avoid-B test pairs ($F[1,51] = 4.7$, $p = .03$). Avoid-B performance was also significantly worse in the on than off medication state

($F[1,51]= 5.0$, $p = .03$). In contrast, medicated patients were numerically, but not significantly, better at choose-A test pairs. The relatively intact choose-A performance in non-medicated patients may reflect an incomplete medication washout. Previously reported detrimental effects on negative learning were also more robust than the beneficial effects on positive learning (*2, 4*). Moreover, although patients off medication appear to be indistinguishable to senior controls, patients required a greater number of trials to reach performance criterion during the training phase before advancing to the critical test phase (see above). This slower probabilistic learning is consistent with other data, and may reflect reduced dynamic range of dopaminergic reinforcement signals (*6, 8*).

Notably, DBS did not alter patients' biases to learn from positive or negative feedback ($p > 0.1$, Figure 1 main paper). Overall, DBS patients performed numerically but not significantly worse than the others at novel test pairs ($F[1,51] = 2.3$, ns). This numerical accuracy deficit is consistent with the overall advanced disease progression in patients treated with DBS. Moreover, there were significant RT differences across conflict conditions (even in correct trials), indicating that learning was sufficient to induce slowing (or speeding) as a function of stimulus-reinforcement conflict, in opposite directions in patients on and off DBS.

Note that classic models of the BG would predict that if the indirect "NoGo" pathway is involved in learning from negative feedback, then DBS should affect this learning, since the STN is part of the indirect pathway (*9, e.g.,*). However, our model suggests that NoGo learning for suppressing the execution of a *specific* response (eg, avoid selecting stimulus B), is implemented via focused projections from striatal NoGo cells to the external pallidum (GPe), and then from GPe via focused projections to the internal pallidum (GPi) (*8, 10*). In contrast, STN projections to the pallidum are diffuse, and are not involved in the learned suppression of an individual response. This is further supported by our simulations (Figure 4b of main paper) showing no effect of STN manipulation on learning biases, consistent with the empirical evidence.

2.4 Accuracy in High-Conflict Test Choices

The model predicts that the STN "Hold your horses" signal is adaptive, and prevents premature responding during high conflict decisions (*10*). Indeed, models with simulated STN lesions fail to slow RTs for high conflict choices, and as a result are less accurate than intact (non-PD) networks at discriminating between subtly different reinforcement values. STN-DBS did not have this direct effect in reducing accuracy on high vs low conflict accuracy in human PD patients. While there was a main effect of conflict on accuracy ($F[1,51]$, $=6.1$, $p = .017$), there was no group by conflict interaction ($F[4,51] = 0.9$), and planned comparisons revealed no difference between patients on vs off DBS ($F[1,51] = 0.14$). However, this null effect is equivocal for multiple reasons. First, overall accuracy on high conflict test pairs was quite low for both patients off and on DBS (57% in both cases, in contrast to patients on and off medication: 68 and 65% respectively, and senior controls: 64%). This overall poor performance in high conflict test choices in patients both on and off DBS makes it difficult to detect a specific high-conflict accuracy deficit in patients on DBS. Further, the low performance in off DBS patients should not be surprising, given that it is difficult to discriminate between slight differences in reinforcement values, and that these patients had relatively advanced Parkinson's disease (Table 1).

Nevertheless, our model makes a clear and stronger prediction that error trials in patients on DBS should specifically be associated with premature responding in high conflict trials, whereas

errors in other groups would stem from other factors. Indeed, as reported in the main paper, analysis of error trials revealed that patients on DBS were still significantly and reliably *faster* for high than low conflict decisions ($F[1,51] = 16.1$, $p = .0002$), supporting the hypothesis that these errors reflected premature responding. In contrast, no such effect was seen in patients off DBS ($F[1,51] = 0.4$), so that their high-conflict errors could not be attributed to premature responding. The same logic holds true for patients on and off medication (p's > 0.15). Healthy senior control subjects even showed the reverse pattern, still showing significantly slower RT's for high than low conflict error trials ($F[1,51] = 9.9$, $p = .003$). This pattern suggests that these seniors adaptively slowed down under high conflict conditions, but that other factors (eg difficulty resolving differences in reinforcement probabilities) led to errors.

Furthermore, we also found that across patients on DBS, the more they showed premature responding, the more they made high conflict errors. This was true regardless of whether we defined premature responding as faster error than correct trials only in high conflict conditions ($r(13) = 0.53$, $p = .05$), or whether we computed a relative measure of premature responding in high compared with low conflict trials ($r(13) = 0.61$, $p = .03$). (One patient did not make any low conflict errors and another patient did not make it to the test phase. Thus this analysis includes the remaining 13 patients.) In sum, although patients on DBS did not perform worse than those off DBS at high conflict test trials, only the high number of errors in the former group could be attributed to premature responding.

Our model also provides a plausible account for why accuracy should be relatively spared under DBS in high conflict trials. Specifically, although simulated STN lesions lead to impaired high-conflict accuracy (*10*), an effect that we replicated here, surprisingly, models with simulated external DBS did not show this accuracy deficit. This discrepancy is explained by the following logic. In high conflict trials, response times have to be slow enough so that the striatal system can discriminate between competing Go signals so as to select the best one. In STN lesioned models, reaction times are significantly speeded, an effect that is exaggerated for win/win decisions, leading to premature responding and elevated errors (50 +/- 8% errors compared with 16 +/- 6% in intact networks). In the DBS networks, although responses were still speeded relative to low conflict choices, they were nevertheless slower than those of the STN-lesioned models, due to overall more STN activation (see Figure 5 in main paper, compare with zero STN activation). The relatively greater time taken by these networks allowed them to choose appropriately even in high conflict conditions (28 +/- 7% errors). These modeling and empirical results suggest that high frequency regular STN firing associated with DBS may actually be somewhat preventive of impulsive decisions, compared with a full STN lesion.

2.5 Lose-Lose Conflict

In the main paper we reported that PD patients on DBS did not only fail to slow reaction times for high conflict choices, but actually showed *speeded* responses for win/win decisions. This pattern supports the notion that reaction times to win/win decisions are normally comprised of two competing factors. First, the presence of two positive stimuli should lead to enhanced stimulus-evoked dopamine release, which can drive speeded response times (*11*). In opposition to this factor, the healthy system can adaptively slow decision times in proportion to conflict (*10*). This depiction predicts that when faced with high conflict lose/lose decisions, only one factor is at

play (the conflict-induced slowing), because increased dopamine should not be seen in this case. Therefore, patients on DBS should not actually speed up responding, but they should also not slow responding. This is the pattern that we observed (Figure 2a). Across all participants there was a significant conflict-induced slowing effect for lose/lose decisions ($F[1,51] = 8.3$, $p = .006$). This pattern was not reversed in patients on DBS as it was for win/win decisions. There was no effect of conflict on RT in this case (the numerical trend for a slowing effect was not close to significant; $F[1,51] = 0.08$). There was also a trend for an interaction between conflict and win/win vs lose/lose conditions in DBS patients ($p=. 09$).

Notably, the same lack of speeded responses was observed in the model (Figure 2b) when DA levels were not elevated (as should be the case for lose/lose decisions) during response selection in the test phase.

2.6 Reaction Time Differences in Medicated vs Non-medicated Patients

Overall, there was a trend for medicated patients to have slower RT's than non-medicated patients, but this effect was not significant ($F[1,51]= 3.2$, $p = .08$). Nevertheless this result is surprising, given that slower RT's are typically observed in non-medicated patients (indeed this is the pattern seen in the simulated Parkinson's model). However, first we note that the critical comparisons are the within-subject effects of low versus high conflict, and both non-medicated and medicated patients showed evidence for conflict-induced slowing, with no differences between the groups. Second, we speculate that the somewhat faster RT's in non-medicated patients may reflect enhanced motivation to perform well in cognitive tasks, and that this motivation can itself speed RTs. These patients are aware that their cognition is being evaluated, and are typically sensitive about this, particularly in the off-medication condition in which they do not have the benefits of pharmacological enhancement. We previously made a similar argument for why these patients can actually perform better than controls in some cases (the so-called Hawthorne effect; Frank et al, 2004).

2.7 STN-DBS and Post-Error Slowing

We also found evidence for the same putative "hold-your-horses" mechanism in the learning phase of the task. Specifically, in our probabilistic learning task, after receiving negative feedback for a given trial (e.g., CD), healthy individuals typically slow reaction times the next time they encounter the same trial type. Across all participants, reaction times were slower for trials that had been most recently associated with negative compared with positive feedback ($F[1,51] = 5.1$, $p = .03$). This phenomenon of post-error slowing (*12*) is thought to reflect the same neural mechanism in the anterior cingulate cortex that detects conflict and leads to cautious behavior (*13, 14*). Participants could potentially leverage this slowing to improve accuracy on the subsequent trial. We reasoned that slower reaction times following negative feedback would be associated with more accurate choices. We therefore hypothesized that this effect would be intact in patients off DBS, but that STN-DBS would interfere with cingulate-STN connectivity and abolish the coupling between post-error reaction time and accuracy.

While there was no overall DBS effect on response time ($F[1,15] = 0.02$), the relative RT slowing following incorrect relative to correct trials marginally depended on DBS status (on vs

off; F[1,15] = 3.9, p = .067). Notably, the degree of post-error slowing was tightly coupled with post-error accuracy in patients off DBS (Figure 3a; r(12) = 0.72, p = .006), whereas this relationship was abolished in patients on DBS (Figure 3b; r(15) = -.03).

2.8 Second Test Phase Effects

Because DBS patients performed both the learning and test phases on DBS, it is theoretically possible that DBS fundamentally altered the manner in which reinforcement contingencies were acquired. That is, although there were no effects of DBS on training accuracy, there were nevertheless other post-error slowing effects and possible differences in the manner in which reinforcement values were acquired, potentially leading to only an indirect conflict effect. The "retrograde DBS" effects reported in the main paper control for this potential confound by showing that even patients who had learned the task off DBS, and had shown conflict-induced slowing in the test phase, showed the reverse speeding effect when tested again with their stimulators turned back on. Furthermore, to provide a treatment control, patients who learned off medication were also tested in a second test phase after taking their regular dose of levodopa medication (here the delay was 40-60 minutes to ensure the medication would have time to absorb). Despite this longer delay (which may tax memory), these patients continued to show a numerical trend for conflict-induced slowing in the second test phase (F[1,51] = 2.0, p=.16). Critically, there was a significant treatment by conflict interaction (F[1,51]= 6.0, p = .017), such that DBS reversed conflict-induced slowing but medication did not.

3 Model Methods

Simulations that demonstrate how to replicate all the reported effects can be obtained by sending an email to the author at mfrank@u.arizona.edu. For animated video captures of model dynamics during response selection and learning, please see `www.u.arizona.edu/~mfrank/BGmodel_movies.html`

The model is implemented using the Leabra framework (*15, 16*). Leabra uses point neurons with excitatory, inhibitory, and leak conductances contributing to an integrated membrane potential, which is then thresholded and transformed via an $x/(x+1)$ sigmoidal function to produce a rate code output communicated to other units (discrete spiking can also be used, but produces noisier results). Synaptic connection weights were trained using a reinforcement learning version of Leabra (*8, 15*). The learning algorithm involves two phases, and is more biologically plausible than standard error backpropagation. In the *minus phase*, the network settles into activity states based on input stimuli and its synaptic weights, ultimately "choosing" a response. In the *plus phase*, the network resettles in the same manner, with the only difference being a change in simulated dopamine: an increase of SNc unit firing from 0.5 to 1.0 for correct responses, and a decrease to zero SNc firing for incorrect responses (*8*). Further model details, including equations and parameters are as described in Frank (2006).

As shown in the main paper, the observed behavioral pattern of results was predicted by our computational model (*8,10*). To explicitly compare these results to the model, we simulated probabilistic learning under multiple model conditions (intact, simulated PD, simulated DA medication,

simulated STN lesions, and simulated external DBS, as described in the main paper).

The network's task was to select one of two possible responses for each stimulus input cue. "Feedback" is then provided to the network by either increasing or decreasing simulated dopamine levels in the plus phase, as described above. The network learns based on the difference in Go/NoGo activity levels in the response selection and feedback phase, as detailed in (*10*).

The stimulus-response mappings are probabilistic, such that the optimal response for some cues will lead to positive reinforcement (DA bursts) in 80% of trials; in the remaining 20% of trials some alternative response is reinforced. For all incorrect responses, DA dips are applied. Other cue-response mappings are less reliable, such that the optimal response is positively reinforced in only 60% of trials. 25 networks with different sets of random initial synaptic weights were run with each manipulation. Each network was trained for 25 epochs (the qualitative pattern of results does not depend on this number), consisting of 10 trials of each stimulus cue.

To determine whether the STN is beneficial for selecting among multiple competing responses, a test phase was administered. Two cues were presented in the input simultaneously, one of which had been associated with 80% positive reinforcement for a given response, while the other had been associated with 60% positive reinforcement for an alternative response. Although the models had not been trained with these stimulus combinations, they should nevertheless be able to select the response that was most likely to result in positive reinforcement (i.e. the 80% response). However, premature responding could result in selection of the alternative suboptimal response if its corresponding striatal Go signal happened to get quickly active (due to noise in striatum or in the premotor representations themselves). This is precisely this kind of situation that an initial STN Global NoGo signal may be useful, so that the network can integrate over multiple possible responses before selecting the most appropriate one.

3.1 Reaction time measures under low and high conflict

To measure reaction times in the model, we assessed the number of network processing cycles before a response was selected by the BG action selection network. We therefore counted a response as being gated by the BG when one of the thalamus units was disinhibited so that its activity exceeded 0.5 (50% maximal firing). These same methods were employed in previous model reaction time analyses (*17*). Because the BG gating system is required to facilitate a cortical response, similar results are obtained by probing Output unit activity.

To simulate win/win decisions and response speeding, DA levels were somewhat elevated during response selection (firing rates of DA neurons were increased from 50% maximal to 70%). This change was applied for both low and high conflict conditions, as long as a single positive stimulus was presented (ie DA levels were not further increased for two vs. one positive stimuli, consistent with recent data (*18*). Decision times were nevertheless faster for win/win decisions, due to the presence of multiple Go signals in the striatum). This relationship between reaction times and DA levels is consistent with that observed in experimental monkeys (*11*). For lose/lose simulations, DA levels were not elevated.

3.2 External DBS Simulation

To simulate DBS, we applied continuous external excitatory input stimulation to a subset of STN units (3 out of 9 units). We also increased the time constant at which STN units could update their membrane potentials (from .05 to 1.0), allowing STN units to oscillate at very high frequency. The resulting oscillations occur at much higher frequency than the slow oscillations reported previously for DA-depleted networks, putatively associated with Parkinson's tremor (*10*). Note that results reported in the DBS simulations do not depend on the particular parameters used to simulate these oscillations – indeed similar results were obtained in yet a third simulation in which we simply added Gaussian noise to each model STN unit's activity. Thus the key finding is that disruption of natural STN processing by externally stimulating, adding noise, or removing the STN altogether will prevent the system from regulating decisions times in proportion to decision conflict.

3.3 Representation of Conflict: Cingulate or PreSMA?

Several studies suggest that the anterior cingulate cortex (ACC) is responsible for detecting response conflict (*19, 20*). These authors suggest that the ACC corresponds to the monkey rostral cingulate motor zone (*21, 22*), and may represent the activation of multiple competing responses. In models of this system, the ACC detects response conflict in the form of the co-activation of multiple motor plans, but integrates this conflict *across time* for subsequent modification of behavioral control (in future trials) (*23, 24*). In contrast, the conflict signal in our model effectively represents *instantaneous* conflict, as the preSMA integrates activity from its sensory inputs within the course of a single trial. Greater co-activation among multiple responses leads to increased conflict output from preSMA to STN. For high conflict trials, the preSMA has increased activity due to the presence of two stimulus inputs, which are both trying to drive activity in the different response units. In future work, it would be interesting to further incorporate an ACC layer that integrates preSMA conflict over time, and interacts with the STN to control behavior between trials.

3.4 Go/NoGo model biases

Following probabilistic training, Go/NoGo associations were recorded from the model's striatum, by computing activation-based receptive fields. For positive Go learning, we computed the summed Go - NoGo activity of units representing the response leading to the selection of the 80% rewarded response. For negative NoGo learning, we computed summed NoGo - Go associations for selecting the alternative 80% negative response (*2*). The Go/NoGo learning results shown in main paper replicate those shown previously, before the STN was included in our model. The STN lesioned from the current version of the model resembles that of the intact network, with respect to Go/NoGo learning biases, and simulated DBS also does not alter this bias.

The tendency for simulated PD networks to learn less from positive outcomes is consistent with previous observations in non-medicated patients across two different tasks (*2*). This effect was numerically but not significantly observed in non-medicated compared with medicated patients in the current sample. However, non-medicated patients took somewhat longer to reach criterion in the training phase (see above), and the residual transfer to positive test outcomes may also reflect an incomplete washout of medication. Previously reported effects on negative learning were also

more robust than those on positive learning (*2, 4*).

References

1. J. R. Blair, O. Spreen, *Clinical Neuropsychologist* **3**, 129 (1989).
2. M. J. Frank, L. C. Seeberger, R. C. O'Reilly, *Science* **306**, 1940 (2004).
3. M. J. Frank, R. C. O'Reilly, *Behavioral Neuroscience* **120**, 497 (2006).
4. R. Cools, L. Altamirano, M. D'Esposito, *Neuropsychologia* **44**, 1663 (2006).
5. C. M. Judd, G. H. McClelland, *Data Analysis, A Model-Comparison Approach* (Harcourt Brace Jovanovich, Orlando, FL, 1989).
6. B. J. Knowlton, J. A. Mangels, L. R. Squire, *Science* **273**, 1399 (1996).
7. D. Shohamy, *et al.*, *Brain* **127**, 851 (2004).
8. M. J. Frank, *Journal of Cognitive Neuroscience* **17**, 51 (2005).
9. R. L. Albin, A. B. Young, J. B. Penney, *Trends in Neurosciences* **12**, 366 (1989).
10. M. J. Frank, *Neural Networks* **19**, 1120 (2006).
11. T. Satoh, S. Nakai, T. Sato, M. Kimura, *Journal of Neuroscience* **23**, 9913 (2003).
12. D. Laming, *Acta Psychologica* **43**, 199 (1979).
13. G. Hajcak, N. McDonald, R. F. Simons, *Psychophysiology* **40**, 895 (2003).
14. M. J. Frank, B. S. Woroch, T. Curran, *Neuron* **47**, 495 (2005).
15. R. C. O'Reilly, Y. Munakata, *Computational Explorations in Cognitive Neuroscience: Understanding the Mind by Simulating the Brain* (The MIT Press, Cambridge, MA, 2000).
16. R. C. O'Reilly, *Neural Computation* **13**, 1199 (2001).
17. M. J. Frank, A. Scheres, S. J. Sherman, *Philosophical Transactions of the Royal Society-B* **362**, 1641 (2007).
18. G. Morris, A. Nevet, D. Arkadir, E. Vaadia, H. Bergman, *Nat Neurosci* **9**, 1057 (2006).
19. M. Botvinick, L. E. Nystrom, K. Fissel, C. S. Carter, J. D. Cohen, *Nature* **402**, 179 (1999).
20. M. M. Botvinick, J. D. Cohen, C. S. Carter, *Trends in Cognitive Sciences* **8**, 539 (2004).
21. N. Picard, P. L. Strick, *Cerebral Cortex* **6**, 342 (1996).
22. N. Picard, P. L. Strick, *Curr Opin Neurobiol* **11**, 663 (2001).

23. M. M. Botvinick, T. S. Braver, D. M. Barch, C. S. Carter, J. D. Cohen, *Psychological Review* **108**, 624 (2001).

24. S. McClure, M. Gilzenrat, J. Cohen, *Advances in Neural Information Processing Systems 18* (2005).

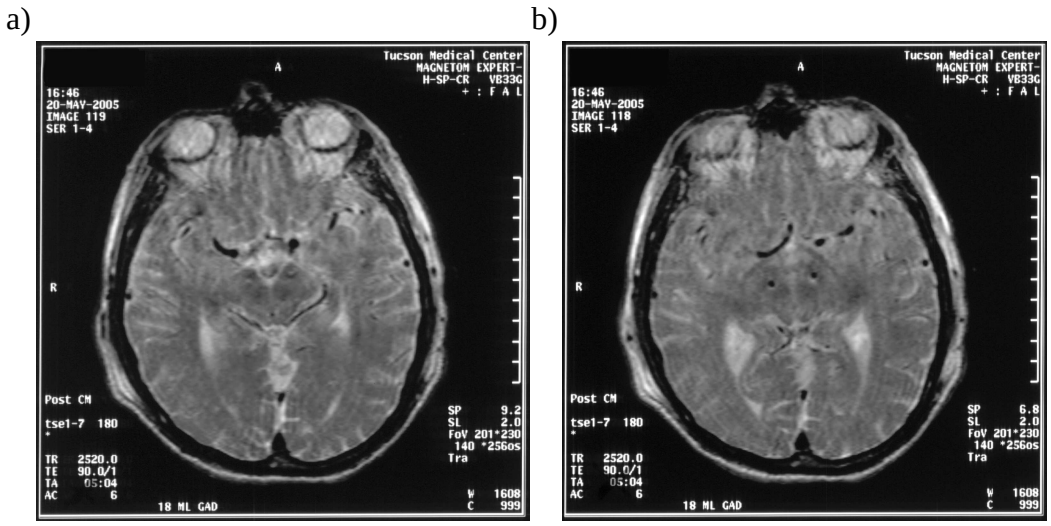

Figure 1: MR images of the subthalamic region **a)** pre-operatively, and **b)** post-operatively. Electrode contacts can be observed in the STN region in (b).

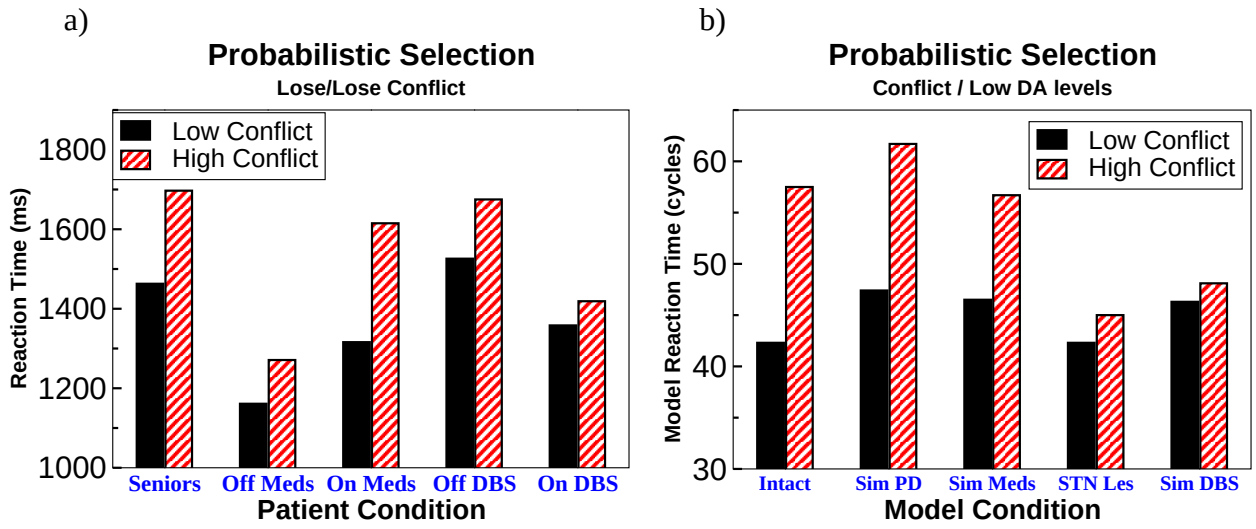

Figure 2: **a)** Mean of median reaction times to high-conflict lose/lose decisions (correct trials) **b)** Model data showing similar pattern of results. In these simulations, model DA levels were not elevated (no positive stimuli), and models with simulated STN lesions or external DBS did not show high-conflict speeding, but still showed reduced slowing than the other cases.

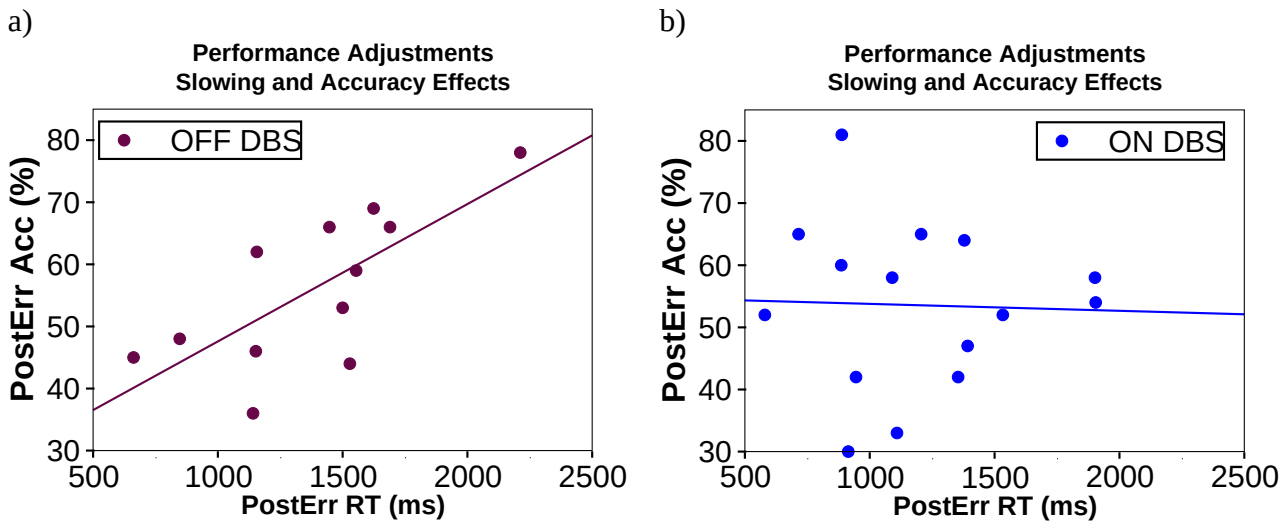

Figure 3: Coupling between trial-specific post-error slowing and post-error accuracy in patients **a)** off and **b)** on STN-DBS.

Group	n	n filt	Sex ratio (m:f)	Age	Years Education	NAART (# correct)	Hoehn & Yahr stage	Years diag
Seniors	27	22	7:15	66.0 (1.7)	16.2 (0.7)	44.0 (1.9)	N/A	N/A
PD patients								
ON Med	15	12	7:5	67.8 (2.1)	17.8 (1.2)	42.9 (2.3)	2.4 (0.2)	8.8 (0.8)
OFF Med	14	9	6:3	67.6 (2.5)	19.2 (1.4)	43.5 (3.1)	2.3 (1.9)	9.5 (1.4)
ON DBS	17	15	13:2	64.5 (2.8)	14.2 (1.5)	39.9 (2.4)	2.3 (0.2)	14.4 (1.5)
OFF DBS	14	12	11:1	62.3 (3.3)	14.4 (1.2)	39.0 (2.9)	2.8 (0.3)	15.2 (1.8)

Table 1: Demographic variables for seniors and PD patients, with no significant differences between groups in any of the demographic variables, except DBS patients having had the disease longer than medication patients (p = .004), as is expected. There were trends for on/off DBS patients to be less educated (p= .08) and more advanced disease progression (p=.13) than patients on/off medication. These patients were not significantly less educated than the senior controls. Groups were not gender-matched, but it is unlikely that this factor impacts on the results given that our DBS and medication manipulations were within-subject. The "n filt" column shows the number of remaining participants who were not filtered out for data analysis (see Data Filtering sections); participants who were filtered out were not included in the demographic means displayed for that row (as they were not used in the statistical comparisons). NAART = number correct responses in the North American Adult Reading Test, an estimate of premorbid verbal IQ. For PD patients, disease severity is indicated in terms of mean Hoehn and Yahr stage, and the number of years since having been diagnosed with PD. Values represent mean (standard error).

Volt-R (V)	Volt-L (V)	Freq-R (Hz)	Freq-L (Hz)	Wave-R (μs)	Wave-L (μs)
3.4 (0.1)	3.2 (0.15)	176.7 (2.9)	170.0 (4.2)	74.0 (3.9)	76.9 (4.7)

Table 2: DBS stimulation parameters for right and left STN. Values represent mean (standard error).