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

Michael J. Frank, *et al.*

Science **318**, 1309 (2007);
DOI: 10.1126/science.1146157

Science (print ISSN 0036-8075; online ISSN 1095-9203) is published weekly, except the last week in December, by the American Association for the Advancement of Science, 1200 New York Avenue NW, Washington, DC 20005.

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

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

Deep brain stimulation (DBS) of the subthalamic nucleus markedly improves the motor symptoms of Parkinson's disease, but causes cognitive side effects such as impulsivity. We showed that DBS selectively interferes with the normal ability to slow down when faced with decision conflict. While on DBS, patients actually sped up their decisions under high-conflict conditions. This form of impulsivity was not affected by dopaminergic medication status. Instead, medication impaired patients' ability to learn from negative decision outcomes. These findings implicate independent mechanisms leading to impulsivity in treated Parkinson's patients and were predicted by a single neurocomputational model of the basal ganglia.

Should you vacation in Montreal or Rome, eat chocolate fondue or tiramisu, go skiing, or visit world-class museums? Such win/win decisions are notoriously difficult to make, often leading to seemingly counterproductive deliberation and hesitation. Intuitively, either option should produce satisfactory results, so why wait? Mathematical models of decision-making suggest that individuals only execute a choice once the "evidence" in its favor crosses a critical decision threshold (*1*, *2*). But the notion of decision threshold need not imply some fixed value. Indeed, individuals can optimally adjust decision thresholds to meet current task demands (*3*–*5*). At the neurobiological level, one model posits that the subthalamic nucleus (STN) dynamically modulates decision thresholds in proportion to reinforcement and decision conflict (*6*). In essence, this model predicts that when faced with multiple seemingly good options, the STN enables you to adaptively "hold your horses," buying more time to settle on the best one. Supporting this account, STN dysfunction in rats causes premature responding in choice paradigms (*7*, *8*). Here we provide direct evidence in humans and show that STN disruption causes impulsive responding during high-conflict win/win decisions.

We administered computerized decision-making tasks to two groups of patients with Parkinson's disease (PD), and age-matched control participants (table S1). One group of patients (n = 17) was tested in different sessions on and off deep brain stimulation (DBS) of the STN, an increasingly common surgical procedure to treat motor symptoms of the disease (*9*). [See (*10*) for DBS surgical procedures, stimulation parameters (table S2), and confirmation of electrode implants in the STN (fig S1).] DBS patients were on relatively low doses of medication in both sessions (*9*, *10*). The second patient group (n = 15) was tested on and off dopaminergic medication. The purpose of the medication manipulation was twofold: (i) to test whether any effects of DBS on conflict-based decisions were selective to that treatment; and (ii) to replicate findings that dopaminergic medication impairs patients' ability to learn from the negative outcomes of their decisions (*11*–*13*), which could account for pathological gambling behavior (*14*). Thus, we tested two potentially distinct mechanisms by which Parkinsonian treatments can cause impulsive behavior (*14*, *15*).

Participants were tested with a probabilistic selection task (*12*, *16*). Three different stimulus pairs (AB, CD, EF) were presented in random order, and participants had to choose one of the two stimuli (Fig. 1A). Feedback followed the choice to indicate whether it was correct or incorrect, but this feedback was probabilistic. In the most reliable AB pair, a choice of stimulus A led to positive feedback in 80% of trials, whereas a choice of stimulus B led to negative feedback in these trials. In a subsequent "test phase," participants chose between novel combinations of all stimuli. Positive-feedback learning was indicated by reliable choice of the most positive stimulus A in these novel test pairs, whereas negative-feedback learning was indicated by reliable avoidance of stimulus B (*12*, *16*). To examine conflict effects, we measured reaction times for test pairs having similar reinforcement values (e.g., 80 versus 70%; high conflict) compared with low-conflict pairs having more easily discriminable values (*6*, *16*) (Fig. 1A).

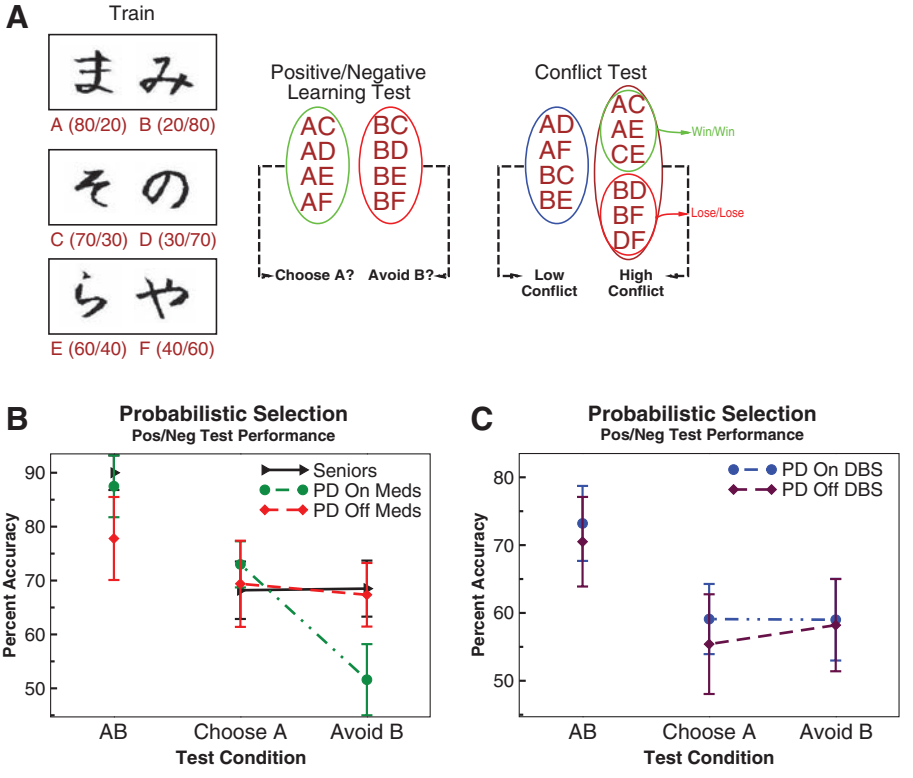

Fig. 1. (**A**) Probabilistic selection task. Each stimulus pair is presented separately in different trials, in random order. Correct choices are determined probabilistically (percent positive/negative feedback is shown in parentheses for each stimulus). A test phase ensues that presents all novel recombinations to assess positive/negative learning biases and conflict (*12*, *16*). (**B**) PD and medication effects, showing selectively impaired avoid-B performance in medicated patients. Nonmedicated patients performed similarly to controls, but were slower to acquire probabilistic contingencies in the learning phase (*10*). (**C**) DBS effects. DBS patients were more advanced in their disease progression than the medication group (table S1); within-patient treatment effects are therefore more interpretable than between-group effects. Error bars are SEs.

[1]Department of Psychology and Program in Neuroscience, University of Arizona, Tucson, AZ 85721, USA. [2]Banner Good Samaritan Medical Center, Phoenix, AZ 85004, USA. [3]Department of Neurology, University of Arizona, Tucson, AZ 85724, USA.

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

It is precisely in these high-conflict choices that it may be adaptive to "hold your horses," increasing the likelihood of settling on the more optimal choice (*6*). We predicted that compared with controls, PD patients (regardless of treatment) would show reinforcement learning deficits (*11*). We further predicted that dopaminergic medication would impair negative-feedback learning (*12*, *13*), whereas DBS would cause impulsive responding in the face of conflict (*6*).

Patients were slower than controls to learn probabilistic reinforcement contingencies (*10*). As shown previously (*12*), patients on medication were selectively impaired at learning from negative decision outcomes [Fig. 1B; see (*10*) for detailed analysis]. Notably, DBS status (on versus off) did not affect positive- or negative-feedback learning (Fig. 1C). Rather, DBS induced impulsive responding under high-conflict conditions. Overall, participants significantly slowed responses for correct high- relative to low-conflict decisions ($F[1,51] = 13.5$, $P < 0.001$; Fig. 2, A and C). This conflict-induced slowing is reminiscent of the deferred decisions under conflict observed in other contexts, including economic decisions (*17*). In contrast, patients on DBS failed to slow down with increased decision conflict (group by conflict interaction, $F[4,51] = 4.9$, $P = 0.002$). The within-subject effect of DBS (on versus off) on conflict-induced slowing was significant ($F[1,51] = 4.6$, $P = 0.036$). Patients on DBS even responded marginally faster under high- than under low-conflict conditions ($F[1,51] = 3.6$, $P = 0.06$). Finally, dopaminergic medication had no effect on conflict-induced slowing ($F[1,51] = 0.5$), and there were no other group/conflict differences.

In models and animals with STN dysfunction, premature responding is associated with suboptimal choices (*6*–*8*). Notably, the tendency for DBS patients to show speeded high-conflict responses was especially pronounced when choosing the less optimal stimulus ("error trials"; $F[1,51] = 16.1$, $P = 0.0002$; Fig. 2, B and D). Further, the more DBS patients exhibited high-conflict premature responding (as defined by faster error than correct choices), the more errors they made in high- than in low-conflict conditions [$r(13) = 0.53$, $P = 0.05$; P's > 0.3 for all other groups]. Thus, high-conflict premature responding led to suboptimal choices under DBS.

Why should DBS patients respond even faster to high- than to low-conflict choices? We posited that the presence of two positive stimuli in high-conflict "win/win" choices could lead to such impulsive responding. Indeed, patients on DBS responded significantly faster during high-conflict win/win conditions (Fig. 3A; $F[1,51] = 5.2$, $P = 0.027$); this faster responding was not observed for lose/lose decisions (fig. S2).

Finally, to control for a possible confounding effect of DBS during the learning phase, we also used a "retrograde DBS" procedure. All patients who had learned the task off DBS were subsequently tested again in a second test phase, identical to the first, after having their stimulators turned on (and a 10-min delay). If DBS genuinely and primarily interferes with the ability to modulate decision times as a function of conflict, these patients should no longer show a conflict-induced slowing effect in the second test phase.

Indeed, the conflict-induced slowing effect was reversed in the second test phase, with DBS patients responding significantly faster to win/win decisions (Fig. 3B; $F[1,51] = 4.7$, $P = 0.03$). These same patients had exhibited the opposite pattern—slowing responses for win/win conditions—just minutes before in the off-DBS state (Fig. 3A). A subset of senior controls ($n = 14$) who were also tested in a second test phase, with the same temporal delay between phases, continued to show conflict-induced slowing ($F[1,51] = 4.3$, $P = 0.04$). 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 [but with a longer delay to allow medication to be absorbed (*10*)]. Critically, there was a significant treatment by conflict interaction ($F[1,51] = 6.0$, $P = 0.017$], such that DBS reversed conflict-induced slowing but medication did not.

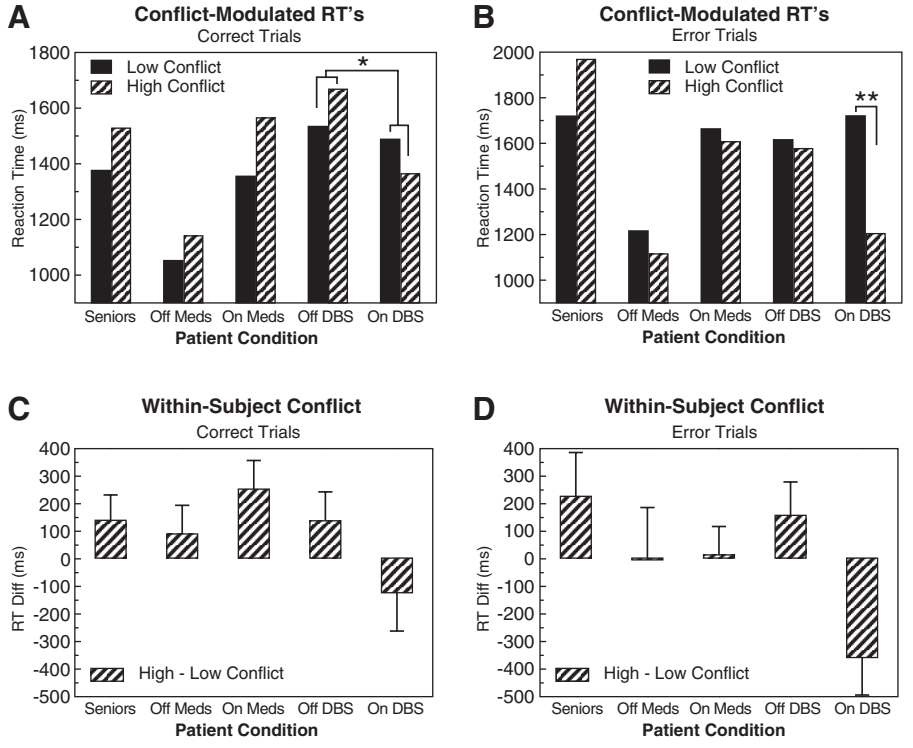

Fig. 2. Conflict effects on decision times. Mean of median reaction times (RTs) are shown for low- and high-conflict conditions in (**A**) correct and (**B**) error trials. Within-subject RT differences (high–low conflict) are also shown in (**C**) correct and (**D**) error trials. The DBS (on versus off) effect on conflict-induced slowing was significant (*$P < 0.05$). Patients on DBS actually responded more rapidly to high-conflict choices, particularly in error trials (**$P < 0.001$).

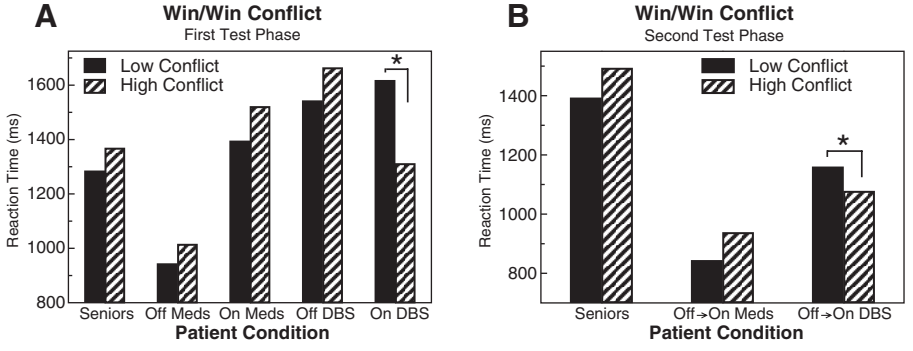

Fig. 3. High-conflict win/win decisions (correct trials; similar results in error trials, not shown). (**A**) Patients on DBS responded significantly faster during high-conflict win/win decisions. (**B**) "Retrograde DBS." Patients who acquired the reinforcement contingencies off DBS were then tested again in a second test phase, after their DBS units were turned back on (*$P < 0.05$).

Taken together, our findings provide evidence for two distinct computational roles of the basal ganglia in decision-making. Dopaminergic medication altered patients' relative tendency to learn from positive versus negative outcomes (*12*, *13*, *18*), without affecting conflict-induced slowing. In contrast, DBS induced speeded high-conflict choices, without affecting learning biases. Both of these findings are captured by a single a priori computational model of the basal ganglia in learning and decision-making (*6*).

As in other models, the basal ganglia in our model supports adaptive decision-making by modulating the selection of frontal cortical action plans (*5*, *6*, *11*, *19*–*21*). In brief, two main neuronal populations in the striatum have opposing effects on action selection via output projections through the globus pallidus, thalamus, and back to the cortex (Fig. 4A). Activity in "Go" neurons facilitates the execution of a cortical response, whereas "NoGo" activity suppresses competing responses. Dopamine bursts and dips that occur during positive and negative outcomes (*22*) drive Go learning (via D1 receptors) to seek rewarding actions, and NoGo learning (via D2 receptors) to avoid actions that are nonrewarding (*11*). Complementing this functionality, the STN provides a self-adaptive dynamic control signal that temporarily prevents the execution of any response, depending on decision conflict (*6*). Notably, the STN receives direct projections from the presupplementary motor area (preSMA) and cingulate cortex regions that detect and integrate response conflict (*23*–*25*). In turn, the STN sends a "Global NoGo" signal via diffuse excitatory projections to basal ganglia output nuclei (*19*, *26*) with consequent inhibition of thalamocortical activity. This STN mechanism provides a means to implement cognitive control, by effectively raising decision thresholds in the face of conflict (*6*). Supporting this notion, neuroimaging studies have found that preSMA and STN coactivation is associated with slowed response times under decision conflict (*25*), and STN-DBS reduces coupling between cingulate and basal ganglia output (*27*).

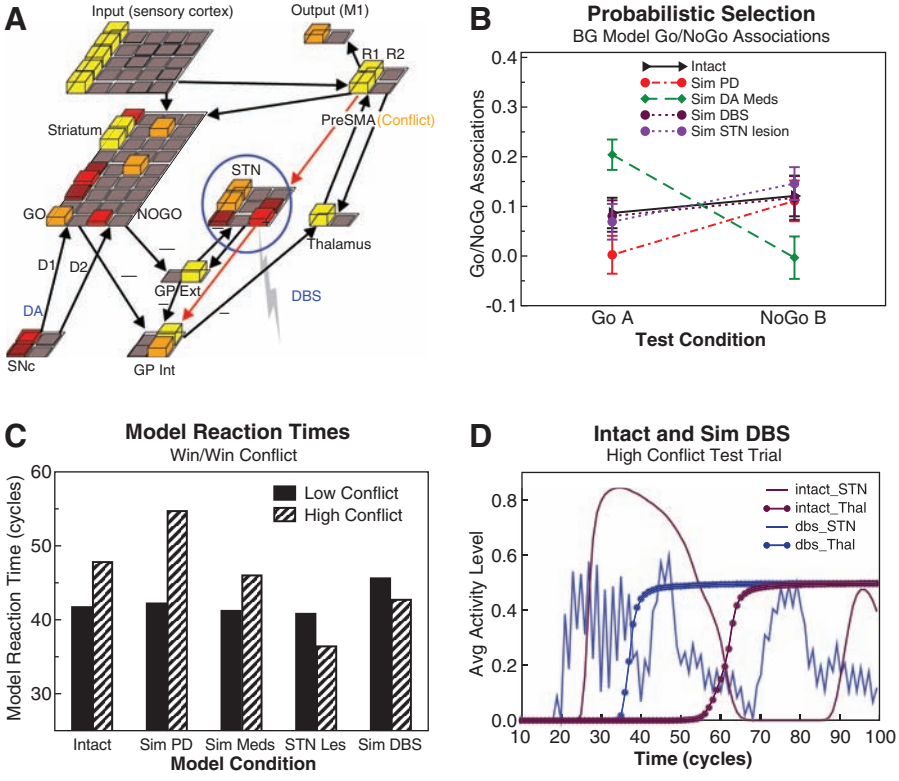

Fig. 4. (**A**) Neural network model of the basal ganglia (squares represent units, with height and color reflecting neural activity). The preSMA selects a response (R1 or R2) via direct projections from the sensory input and is modulated by basal ganglia (BG) output via the thalamus. Go and NoGo units are, respectively, in the left and right halves of the striatum, with separate columns for each response, and receive dopaminergic (DA) learning signals from the substantia nigra pars compacta (SNc). The STN sends a Global NoGo signal by exciting globus pallidus, internal segment (GP Int) in proportion to response conflict in preSMA (these projections shown in red). In the case shown, conflict is low because only a single response (R1) is active. (**B**) Model predictions for reinforcement learning. Plots show striatal activation-based receptive fields indicating summed Go-A and NoGo-B associations (*10*, *12*). (**C**) The same model's predictions for conflict-induced slowing. Reaction times are indexed by the number of processing cycles before a given response is selected (*10*). Simulation results reflect mean values across 25 network runs with random initial synaptic weights. (**D**) Normalized activity in the model STN and thalamus, in a representative high-conflict win/win trial. The model selects a response when thalamus activity rises. The model selects a response when thalamus activity rises and subsequently facilitates the associated preSMA units.

This single model captures both medication and DBS effects, as revealed in computational simulations of these dynamics [Fig. 4, B to D; see (*6*, *10*) for detailed modeling methods]. To simulate PD, we decreased dopamine levels. To simulate medication, we maintained relatively elevated dopamine levels but prevented them from sufficiently decreasing during negative feedback, due to tonic D2 stimulation (*11*, *12*). The resulting Go/NoGo learning and medication effects replicate those reported with an earlier basal ganglia model that did not include the STN (*12*) and show that these learning biases are insensitive to STN manipulation. Here, we focus on DBS simulations in high-conflict decision-making.

The mechanisms underlying the therapeutic effects of DBS are controversial (*28*). One dominant theory is that high-frequency DBS paradoxically acts like a lesion [e.g., via "depolarization block" (*29*)]. Indeed, like DBS, both real and simulated STN lesions ameliorate abnormal network oscillations and motor symptoms of PD (*6*, *30*). To simulate STN lesions, we simply removed STN from processing altogether (*6*). A second theory is that DBS induces regular high-frequency STN firing patterns (*31*) and actually enhances STN output (*28*). To simulate this version, we externally applied high-frequency excitatory input to the STN [see (*10*) for details]. We posited that either DBS mechanism would prevent the STN from naturally and dynamically responding to its cortical inputs and would therefore disrupt conflict-induced slowing.

After training networks with probabilistic reinforcement, we tested them with low- and high-conflict trials, in which two responses were associated with competing reinforcement probabilities (*6*, *10*). Whereas intact networks showed substantial conflict-induced slowing, those with either STN lesions or external stimulation exhibited the same speeded win/win responding observed in DBS patients (Fig. 4C). This "impulsive" speeding resulted from the presence of two striatal "Go" unit populations (one for each rewarding response), which enhanced the probability that one of them surpassed threshold (*6*). Counteracting this factor, in intact networks, cortical response conflict led to an initial STN surge, postponing responding until this STN activity subsided (Fig. 4D). In STN-lesioned networks, there was no such mechanism to allow this slowing to occur. In networks with external stimulation, the idiosyncratic, non–task-related STN firing prevented it from responding naturally and adaptively to conflict signals. Accordingly, the network could select a response earlier in time (Fig. 4D). In sum, our model captures the main observed effects of different PD treatments and in so doing may reveal different computational functions of the basal ganglia in decision-making.

STN dysfunction does not lead to impulsivity in all behavioral situations. For example, STN-lesioned rats show enhanced preference for choices that lead to large delayed rewards compared with those that yield small immediate rewards, in so-called delay-discounting tasks (*32*, *33*). Thus, the STN is not required to value large rewards per se, or even to "wait" for the reward once the response is made. Our model predicts that STN-lesioned rats would indeed respond impulsively in a modified discounting paradigm in which larger rewards could only be obtained by delaying the response itself.

Clinically, our findings point to two mechanisms that lead to distinct forms of impulsivity in treated Parkinson's patients. Dopaminergic medication, by tonically elevating dopamine levels and stimulating D2 receptors, prevents learning from negative decision outcomes (*11*, *13*, *18*). This mechanism may explain pathological gambling behavior in patients treated with D2 agonists (*14*). It is possible that genetic factors conspire with medication to induce impulsive behaviors, given that a gene coding for striatal D2 receptor density predicts negative outcome learning (*34*). In contrast, DBS patients typically take substantially lower doses of medication (*9*), as was the case here (*10*). Their impulsive decision-making (*15*, *35*) may be explained by an inability to self-modulate decision times as a function of conflict. For example, the first DBS patient in our study, when asked whether he might be more comfortable in a different chair situated across the room, immediately advanced toward that chair, ignoring the fact that he was not able to walk properly and was likely to fall. It is plausible that the "rewarding" prospect of the comfortable chair was not appropriately offset by a functional STN that would have prevented him from reacting so rashly. Such anecdotal evidence is supported by laboratory studies showing DBS-induced impairments in cognitive control (*27*, *36*). Thus, future research should evaluate alternative DBS protocols that take cognitive conflict into account.

References and Notes

1. R. Ratcliff, T. Van Zandt, G. McKoon, *Psychol. Rev.* **106**, 261 (1999).
2. M. Usher, J. L. McClelland, *Psychol. Rev.* **108**, 550 (2001).
3. R. Bogacz, E. Brown, J. Moehlis, P. Holmes, J. D. Cohen, *Psychol. Rev.* **113**, 700 (2006).
4. P. Simen, J. D. Cohen, P. Holmes, *Neural Netw.* **19**, 1013 (2006).
5. C.-C. Lo, X.-J. Wang, *Nat. Neurosci.* **9**, 956 (2006).
6. M. J. Frank, *Neural Netw.* **19**, 1120 (2006).
7. C. Baunez, T. W. Robbins, *Eur. J. Neurosci.* **9**, 2086 (1997).
8. C. Baunez, A. Christakou, Y. Chudasama, C. Forni, T. W. Robbins, *Eur. J. Neurosci.* **25**, 1187 (2007).
9. A. L. Benabid, *Curr. Opin. Neurobiol.* **13**, 696 (2003).
10. Supporting material is available on *Science* Online.
11. M. J. Frank, *J. Cognit. Neurosci.* **17**, 51 (2005).
12. M. J. Frank, L. C. Seeberger, R. C. O'Reilly, *Science* **306**, 1940 (2004).
13. R. Cools, L. Altamirano, M. D'Esposito, *Neuropsychologia* **44**, 1663 (2006).
14. M. L. Dodd *et al.*, *Arch. Neurol.* **62**, 1377 (2005).
15. J. A. Saint-Cyr, A. Albanese, *Neurology* **66**, 1799 (2006).
16. M. J. Frank, B. S. Woroch, T. Curran, *Neuron* **47**, 495 (2005).
17. A. Tversky, E. Shafir, *Psychol. Sci.* **3**, 358 (1992).
18. M. J. Frank, R. C. O'Reilly, *Behav. Neurosci.* **120**, 497 (2006).
19. J. W. Mink, *Prog. Neurobiol.* **50**, 381 (1996).
20. M. D. Humphries, R. D. Stewart, K. N. Gurney, *J. Neurosci.* **26**, 12921 (2006).
21. R. Bogacz, K. Gurney, *Neural Comput.* **19**, 442 (2007).
22. W. Schultz, *J. Neurophysiol.* **80**, 1 (1998).
23. M. Botvinick, L. E. Nystrom, K. Fissel, C. S. Carter, J. D. Cohen, *Nature* **402**, 179 (1999).
24. G. Orieux, C. Francois, J. Feger, E. C. Hirsch, *J. Neurosci.* **22**, 8762 (2002).
25. A. R. Aron, T. E. Behrens, S. Smith, M. J. Frank, R. A. Poldrack, *J. Neurosci.* **27**, 3743 (2007).
26. A. Parent, L. Hazrati, *Brain Res. Brain Res. Rev.* **20**, 128 (1995).
27. S. Thobois *et al.*, *J. Cereb. Blood Flow Metab.* **27**, 1173 (2007).
28. Y. Liu, N. Postupna, J. Falkenberg, M. E. Anderson, *Neurosci. Biobehav. Rev.* 10.1016/j.neubiorev.2006.10.007 (2006).
29. A. Benazzouz, M. Hallett, *Neurology* **55**, S13 (2000).
30. H. Bergman, T. Wichmann, M. R. DeLong, *Science* **249**, 1436 (1990).
31. W. Meissner *et al.*, *Brain* **128**, 2372 (2005).
32. C. A. Winstanley, C. Baunez, D. E. H. Theobald, T. W. Robbins, *Eur. J. Neurosci.* **21**, 3107 (2005).
33. J. M. Uslaner, T. E. Robinson, *Eur. J. Neurosci.* **24**, 2345 (2006).
34. M. J. Frank, A. A. Moustafa, H. Haughey, T. Curran, K. Hutchison, *Proc. Natl. Acad. Sci. U.S.A.* **104**, 16311 (2007).
35. S. Drapier *et al.*, *J. Neurol.* **252**, 583 (2005).
36. T. Hershey *et al.*, *Neurology* **62**, 1110 (2004).

37. We thank M. Cohen and M. van't Wout for helpful comments, M. Kim for data collection, and E. Carter for assistance with neural network simulations.

Supporting Online Material

www.sciencemag.org/cgi/content/full/1146157/DC1
Materials and Methods
Figs. S1 to S3
Tables S1 and S2
References

6 June 2007; accepted 15 October 2007
Published online 25 October 2007;
10.1126/science.1146157
Include this information when citing this paper.