By Carrot or by Stick: Cognitive Reinforcement Learning in Parkinsonism

Michael J. Frank,^{1*} Lauren C. Seeberger,² Randall C. O'Reilly^{1*}

¹Department of Psychology and Center for Neuroscience, University of Colorado Boulder, Boulder, CO 80309-0345, USA.

²Colorado Neurological Institute Movement Disorders Center, Englewood, CO 80113, USA.

*To whom correspondence should be addressed. E-mail: frankmj@psych.colorado.edu (M.J.F.); oreilly@psych.colorado.edu (R.C.O.).

To what extent do we learn from the positive versus negative outcomes of our decisions? The neuromodulator dopamine plays a key role in these reinforcement learning processes. Patients with Parkinson's disease, who have depleted dopamine in the basal ganglia, are impaired in tasks that require learning from trial and error. Here we show, using two cognitive procedural learning tasks, that Parkinson's patients off medication are better at learning to avoid choices that lead to negative outcomes than they are at learning from positive outcomes. Dopamine medication reverses this bias, making patients more sensitive to positive than negative outcomes. This pattern was predicted by our biologically-based computational model of basal ganglia/dopamine interactions in cognition, which has separate pathways for "Go" and "NoGo" responses that are differentially modulated by positive and negative reinforcement.

Should you shout at your dog for soiling the carpet, or praise him when he does his business in the yard? Most dog trainers will tell you that the answer is both. The proverbial "carrot and stick" motivational approach refers to the use of a combination of positive and negative reinforcement: one can persuade a donkey to move either by dangling a carrot in front of it or by striking it with a stick. Both carrots and sticks are important for instilling appropriate behaviors in humans. For instance, when mulling over a decision, one considers both pros and cons of various options, which are implicitly influenced by positive and negative outcomes of similar decisions made in the past. Here we report that whether one learns more from positive or negative outcomes varies with alterations in dopamine levels caused by Parkinson's disease and the medications used to treat it.

To better understand how healthy people learn from their decisions (both good and bad), it is instructive to examine under what conditions this learning is degraded. Notably, patients with Parkinson's disease are impaired in cognitive tasks that require learning from positive and negative feedback (1–3). A likely source of these deficits is depleted levels of the neuromodulator dopamine in the basal ganglia of Parkinson's patients (4), because dopamine plays a key role

in reinforcement learning processes in animals (5). A simple prediction of this account is that cognitive performance should improve when patients take medication that elevates their dopamine levels. However, a somewhat puzzling result is that dopamine medication actually worsens performance in some cognitive tasks, despite improving it in others (6, 7).

Computational models of the basal ganglia/dopamine system provide a unified account that reconciles the above pattern of results and makes explicit predictions about the effects of medication on "carrot and stick" learning (8, 9). These models simulate transient changes in dopamine that occur during positive and negative reinforcement, and their differential effects on two separate pathways within the basal ganglia system. Specifically, dopamine is excitatory on the direct or "Go" pathway, which helps facilitate responding, while it is inhibitory on the indirect or "NoGo" pathway, which suppresses responding (10–13). In animals, phasic bursts of dopamine cell firing are observed during positive reinforcement (14, 15) which are thought to act as "teaching signals" that lead to the learning of rewarding behaviors (14, 16). Conversely, choices that do not lead to reward [and aversive events, according to some studies (17)] are associated with dopamine dips that drop below baseline (14, 18). Similar dopamine-dependent processes have been inferred to occur in humans during positive and negative reinforcement (19, 20). In our models, dopamine bursts increase synaptic plasticity in the direct pathway while decreasing it in the indirect pathway (21, 22), supporting "Go" learning to reinforce the good choice. Dips in dopamine have the opposite effect, supporting "NoGo" learning to avoid the bad choice (8, 9).

A central prediction of our models is that non-medicated Parkinson's patients are impaired at learning from positive feedback (bursts of dopamine; "carrots"), due to reduced levels of dopamine. However, the models also make the counterintuitive prediction that patients should display enhanced learning from negative feedback (dips in dopamine; "sticks"), due to their low dopamine levels that facilitate this kind of learning. Conversely, we predict that patients on medication have sufficient dopamine to learn from positive

feedback, but would be relatively impaired at learning from negative feedback due to the medication blocking the effects of normal dopamine dips. This pattern of dopamine effects explains the existing puzzling results in the Parkinson's disease literature showing both cognitive enhancements and impairments from medication (8).

This paper presents a more direct test of the model's predictions. We employed "procedural learning" (i.e., trial and error) tasks (23) with 30 Parkinson's patients and 19 healthy seniors matched for age, education, and an estimate of verbal IQ (see table S1 in (24) for demographic details and N's per task condition). Two different procedural learning tasks were used, one probabilistic and one deterministic, with the task selected at random for the first session. A subset of participants returned for a second session to perform the other task, and Parkinson's patients in this session abstained from taking their regular dose of dopamine medication for a mean of 18 hours prior to the experiment (7, 24).

In the Probabilistic Selection task, three different stimulus pairs (AB, CD, EF) are presented in random order and participants have to learn to choose one of the two stimuli (Fig. 1A). Feedback follows the choice to indicate whether it was correct or incorrect, but this feedback is probabilistic. In AB trials, a choice of stimulus A leads to correct (positive) feedback in 80% of AB trials, whereas a B choice leads to incorrect (negative) feedback in these trials (and vice-versa for the remaining 20% of 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. Note that learning to choose A over B could be accomplished either by learning that choosing A leads to positive feedback, or that choosing B leads to negative feedback (or both). To evaluate whether participants learned more about positive or negative outcomes of their decisions, we subsequently tested them with novel combinations of stimulus pairs involving either an A (AC, AD, AE, AF) or a B (BC, BD, BE, BF); no feedback was provided. We predict that Parkinson's patients on medication, compared with those off medication, learn more from positive than negative feedback, and should therefore reliably choose the best "carrot" (stimulus A) in all novel test pairs in which it is present. In contrast, those off medication should learn more from negative than positive feedback, and therefore reliably avoid the worst "stick" (stimulus B).

In the Implicit Transitive Inference task (25), the reinforcement for each stimulus pair is deterministic, but stimulus pairs are partially overlapping (Fig. 1a). Four pairs of stimuli are presented: A+B−, B+C−, C+D−, and D+E− where + and − indicate positive and negative feedback. A hierarchy (A>B>C>D>E) emerges in which stimuli near the top of the hierarchy develop net positive associative

strengths, while those near the bottom develop net negative associative strengths (25–27). This explains why, when presented with novel combination BD, participants (and animals trained in similar paradigms) often correctly choose stimulus B, despite having no explicit awareness of any hierarchical structure among the items (25, 26, 28, 29). Based on this associative account, we predicted that Parkinson's patients on medication, compared to those off medication, would learn more about the positive associations at the top of the hierarchy, resulting in better performance on stimulus pairs AB and BC. Conversely, those off medication should learn more about the negative associations at the bottom of the hierarchy, resulting in better CD and DE performance. Note that because the novel BD pair could be solved either by a positive B association or a negative D association, we did not predict a difference in BD performance between groups.

Results confirmed our predictions. Despite no main effect of medication, session, or test condition, the critical interaction between medication and test condition was significant for both the probabilistic selection [$F(1,26) = 4.3$, $p < .05$] and transitive inference [$F(1,39) = 5.5$, $p < .05$] tasks. In the probabilistic selection task (Fig. 1b), patients on medication tended to choose stimulus A, indicating that they had found the best carrot in the bunch. In contrast, patients off medication had a greater tendency to avoid stimulus B, indicating that they had learned to avoid the harshest stick. Aged matched controls did not differ in performance between A and B pairs. In the transitive inference task (Fig. 1C), patients on medication performed better at choosing positively associated stimuli at the upper end of the hierarchy, whereas those off medication more reliably avoided negative stimuli at the lower end. Finally, aged matched controls did not differ between performance on pairs at the high and low end of the stimulus hierarchy. There was also no effect of medication on performance on novel pairs AE and BD [$F(1,39) = 1.6$, n.s.].

To compare results across both tasks, we converted accuracy measures for both positive and negative conditions to Z-scores (Fig. 1D). This analysis confirmed a significant interaction between positive/negative condition and Parkinson's disease medication group [$F(1,68) = 10.4$, $p = 0.0019$]. Planned comparisons revealed that patients on medication chose positive stimuli more reliably than they avoided negative stimuli [$F(1,25) = 4.98$, $p < .05$], and more reliably than the other two groups [$F(1,69) = 4.8$, $p < .05$]. Conversely, patients off medication avoided negative stimuli more reliably than they chose positive stimuli [$F(1,15) = 5.42$, $p < .05$], and more reliably than the other two groups [$F(1,69) = 7.6$, $p < 0.05$]. This was also true relative to healthy seniors alone [$F(1,69) = 4.6$, $p < 0.05$].

This last observation is a rare example of enhanced cognitive performance associated with neurological disease,

as it suggests that non-medicated patients made better use of negative feedback than either patients on medication or healthy seniors. Trial-to-trial analysis confirmed that a change of choice behavior in the probabilistic selection task (e.g., they chose C in a CD trial after having chosen D in the previous CD trial), was more often accounted for by negative feedback in the previous trial in patients off medication compared with those on medication [$F(1,26) = 5.62$, $p < .05$]. Medicated patients switched choices just as often during training, but were not as influenced by negative feedback to do so. There was no difference between these groups in the efficacy of positive feedback to modify behavior on a trial-to-trial basis [$F(1,26) = 0.42$, n.s.].

Taken together, these findings provide a mechanistic understanding of the nature of the cognitive sequelae of Parkinson's disease, which ties together a variety of other observations across multiple levels of analysis. First, we build on claims that learning from error feedback is primarily affected in Parkinson's disease (3), by showing that the direction of this effect interacts critically with the valence of the feedback and the medication status of the patient. Second, these results are consistent with neuroimaging studies showing that positive and negative feedback have differential effects on basal ganglia activity (30, 31). Third, they help clarify the basis for why medication sometimes improves but sometimes impairs cognitive deficits in Parkinson's disease, depending on the task (6–8). Specifically, patients on medication displayed enhanced positive feedback learning beyond even that of healthy participants, supporting the idea that medication results in higher than normal amounts of dopamine in ventral striatum, which is relatively spared in early stage Parkinson's disease (4, 6, 7). Finally, our observation that non-medicated patients display enhanced ability to avoid negative stimuli may provide the fundamental basis for reports of enhanced harm avoidance behavior in these patients (32, 33).

An equally important contribution of this work is in its confirmation of very specific predictions made by our computational model of the basal ganglia system (8, 9, 34). Almost all of the basic mechanisms of this model have been postulated in various forms by other researchers. Nevertheless, it represents an integration of these mechanisms into a coherent, mechanistically explicit system. At the most general level, the basal ganglia in our model modulates the selection of actions being considered in frontal cortex (8, 34–36). More specifically, two main projection pathways from the striatum go through different basal ganglia output structures on the way to thalamus and up to cortex (Fig. 2A). Activity in the direct pathway sends a “Go” signal to facilitate the execution of a response considered in cortex, whereas activity in the indirect pathway sends a “NoGo” signal to suppress competing responses. Transient changes in

dopamine levels that occur during positive and negative feedback have opposite effects on D1 and D2 receptors, which are relatively segregated in the direct and indirect pathways, respectively (10–13). Thus the net effects of dopamine bursts during positive reinforcement are to activate the Go pathway and deactivate the NoGo pathway, driving learning so that reinforced responses are subsequently facilitated. Conversely, decreases in dopamine during negative reinforcement have the opposite effect, driving NoGo learning so that incorrect responses are subsequently suppressed/avoided (8).

These dopamine modulation effects on the Go and NoGo pathways lead directly to the predictions that we confirmed in the experiments reported earlier, as revealed in computational simulations of these dynamics (Fig. 2,B and C) (8). To simulate Parkinson's disease, we decreased tonic and phasic dopamine levels in the SNc layer of the network, which reduced the ability to generate dopamine bursts during positive feedback. Therefore, the model was relatively impaired at reinforcing Go firing to correct responses. Furthermore, the low tonic dopamine levels produced a persistent bias on the system in favor of the NoGo pathway, which resulted in a corresponding bias to learn NoGo in response to negative feedback. Thus, in our simulation of the probabilistic selection task (Fig. 2C), the simulated Parkinson's model learned more to NoGo to B instead of Go to A. In contrast, intact models learned an even balance of Go to A and NoGo to B.

To simulate the effects of Parkinson's disease medication, we increased the dopamine levels (both tonic and phasic), but we also decreased the size of the phasic dopamine dips during negative feedback. This latter effect is included because D2 agonist medications taken by the vast majority of our Parkinson's patients (in addition to L-Dopa) tonically bind to D2 receptors irrespective of phasic changes in dopamine firing, thereby “filling in” the dips. The net result is the opposite of our simulated Parkinson's model. The tonic elevation in dopamine receptor activation produced a Go bias in learning, while the diminished phasic dip decreased the model's ability to learn NoGo from negative feedback. These combined effects produced the clear crossover-interaction pattern that we observed in our studies (Fig. 2C). Similar results held for our simulation of the transitive inference experiment (24). Finally, reversal learning deficits observed in Parkinson's patients on medication (6, 7) were also accounted for by this same model (8).

Nevertheless, the model does not capture the overall better performance of Parkinson's patients in our study relative to healthy senior controls. This result is somewhat surprising, given that patient impairments have been observed in previous studies (1–3). One potential reason for this discrepancy is the relative simplicity of our task relative to

those used in previous studies. Furthermore, although our control group was matched to the patients in all of our demographic variables, other uncontrolled variables might have led to differences in overall performance levels. For example, because we had access to patient medical records, we may have successfully excluded more patients than seniors for other age-related neurological impairments. Alternatively, Parkinson's patients may have had greater motivation to perform well, given that they were aware that we were studying cognitive sequelae of their disease (the so-called "Hawthorne Effect"). Further, while abstract neural models can make qualitative predictions (such as the cross-over interactions observed in this study), the quantitative aspects of the predictions require more detailed knowledge of specific parameters of the neural system, along with the precise degree of dopamine depletion and remediation by medication in Parkinson's disease; these data are not available. Therefore, we argue that the most meaningful comparisons are the on-vs-off medication conditions, for which the model and data are in close agreement. In addition, the model accurately predicts that healthy seniors did not differ in their tendency to learn from positive versus negative feedback. Finally, we note that our model does not explicitly consider the uneven levels of dopamine depletion in ventral and dorsal striatum of Parkinson's patients (4), which are also thought to play a role in cognitive enhancements/impairments resulting from medication (6–8).

In summary, we have presented evidence for a mechanistic account of how the human brain implicitly learns to make choices that lead to good outcomes, while avoiding those that lead to bad outcomes. The consistent results across tasks (one probabilistic and the other deterministic), and in both medicated and non-medicated Parkinson's patients, provide substantial support for a dynamic dopamine model of cognitive reinforcement learning.

References and Notes

1. B. J. Knowlton, J. A. Mangels, L. R. Squire, *Science* **273**, 1399 (1996).
2. F. Ashby, L. Alfonso-Reese, A. Turken, E. Waldron, *Psychological Review* **105**, 442 (1998).
3. D. Shohamy, *et al.*, *Brain* **127**, 1 (2004).
4. S. Kish, K. Shannak, O. Hornykiewicz, *New England Journal of Medicine* **318**, 876 (1988).
5. R. A. Wise, P.-P. Rompre, *Annual Review of Psychology* **40**, 191 (1989).
6. R. Swainson, *et al.*, *Neuropsychologia* **38**, 596 (2000).
7. R. Cools, R. Barker, B. Sahakian, T. Robbins, *Cerebral Cortex* **11**, 1136 (2001).
8. M. Frank, *Journal of Cognitive Neuroscience* (in press).
9. R. C. O'Reilly, M. J. Frank, Making working memory work: A computational model of learning in the frontal cortex and basal ganglia, *ICS Tech Report 03-03*, revised

- 8/04, University of Colorado Institute of Cognitive Science (2003).
10. C. Gerfen, *Annual Review of Neuroscience* **15**, 285 (1992).
11. I. Aubert, I. Ghorayeb, E. Normand, B. Bloch, *Journal of Comparative Neurology* **418**, 22 (2000).
12. S. Hernandez-Lopez, J.argas, D. Surmeier, A. Reyes, E. Galarraga, *Journal of Neuroscience* **17**, 3334 (1997).
13. S. Hernandez-Lopez, *et al.*, *Journal of Neuroscience* **20**, 8987 (2000).
14. W. Schultz, *Neuron* **36**, 241 (2002).
15. W. Schultz, P. Dayan, P. R. Montague, *Science* **275**, 1593 (1997).
16. J. Wickens, *Network: Computation in Neural Systems* **8**, R77 (1997).
17. M. Ungless, P. Magill, J. Bolam, *Science* **303**, 2040 (2004).
18. T. Satoh, S. Nakai, T. Sato, M. Kimura, *Journal of Neuroscience* **23**, 9913 (2003).
19. C. B. Holroyd, M. G. H. Coles, *Psychological Review* **109**, 679 (2002).
20. D. Zald, *et al.*, *Journal of Neuroscience* **24**, 4105 (2004).
21. A. Nishi, G. Snyder, P. Greengard, *Journal of Neuroscience* **17**, 8147 (1997).
22. D. Centonze, B. Picconi, P. Gubellini, G. Bernardi, P. Calabresi, *European Journal of Neuroscience* **13**, 1071 (2001).
23. Note that although participants may become aware of which stimulus they are actually choosing, they are likely not aware of the detailed influence of the history of positive and negative feedback on their choices. This is the sense in which we regard these as procedural or implicit learning tasks.
24. Materials and methods are available as supporting online material on *Science Online*.
25. M. Frank, J. Rudy, W. Levy, R. O'Reilly, *Memory and Cognition* (in press).
26. M. Frank, J. Rudy, R. O'Reilly, *Hippocampus* **13**, 341 (2003).
27. L. von Fersen, C. D. L. Wynne, J. D. Delius, J. E. R. Staddon, *Journal of Experimental Psychology: Animal Behavior Processes* **17**, 334 (1991).
28. C. Wynne, *Animal Learning & Behavior* **23**, 207 (1995).
29. J. Delius, M. Siemann, *Behavioural Processes* **42**, 107 (1998).
30. M. R. Delgado, H. M. Locke, V. Stenger, J. A. Fiez, *Cognitive, Affective, and Behavioral Neuroscience* **3**, 27 (2003).
31. C. Zink, G. Pagnoni, M. Martin-Skurski, C. J.C., G. Berns, *Neuron* **42**, 509 (2004).
32. V. Kaasinen, *et al.*, *Proceedings of the National Academy of Sciences* **98**, 13272 (2001).

33. R. Tomer, J. Aharon-Peretz, *Journal of Neurology, Neurosurgery & Psychiatry* **75**, 972 (2004).
34. M. J. Frank, B. Loughry, R. C. O'Reilly, *Cognitive, Affective, and Behavioral Neuroscience* **1**, 137 (2001).
35. J. Mink, *Progress in Neurobiology* **50**, 381 (1996).
36. G. Chevalier, J. M. Deniau, *Trends in Neurosciences* **13**, 277 (1990).
37. Supported by ONR grants N00014-00-1-0246 and N00014-03-1-0428, and NIH grant MH069597-01. We thank the staff of the Colorado Neurological Institute for their assistance, and Karen Richardson for help in administering cognitive tasks to patients and senior participants. We also thank Tim Curran, Gary McClelland, and Yuko Munakata for helpful comments.

Supporting Online Material

www.sciencemag.org/cgi/content/full/1102941/DC1

Materials and Methods

Fig. S1

Tables S1 to S3

References and Notes

19 July 2004; accepted 27 October 2004

Published online 4 November 2004; 10.1126/science.1102941

Include this information when citing this paper.

Fig. 1. a) Example stimulus pairs (Hiragana characters) used in both cognitive procedural learning tasks, designed to minimize verbal encoding. One pair is presented per trial, and the participant makes a forced choice. In Probabilistic Selection, the frequency of positive feedback for each choice is shown. In Transitive Inference, feedback is deterministic and indicated by the +/- signs for each stimulus. Any of 12 keys on the left side of the keyboard selects the stimulus on the left, and vice-versa for the right stimulus. The stimulus locations were randomized across trials, and assignment of Hiragana character to stimulus label (A-F) was randomized across participants. In actuality, different Hiragana characters were used across tasks. **b)** Novel test pair performance in the probabilistic selection task, where choosing A depends on having learned from positive feedback, while avoiding B depends on having learned from negative feedback. **c)** Training pair performance during the test phase in the transitive inference task. Stimuli at the top of the hierarchy (A, B) have net positive associations, while those at the bottom (C, D) have net negative associations (24–29). Thus, learning from positive feedback benefits performance on AB & BC, while learning from negative feedback benefits CD & DE. Groups did not differ in novel test pairs AE and BD [not shown; see (24)] which could be solved either by choosing stimuli with positive associations or avoiding those with negative associations. **d)** Z-scores across both probabilistic selection and transitive inference tasks. Positive and negative

conditions correspond to A & B pairs in the probabilistic selection task, and AB/BC & CD/DE pairs in the transitive inference task. Error bars reflect standard error.

Fig. 2. a) The cortico-striato-thalamo-cortical loops, including the direct (“Go”) and indirect (“NoGo”) pathways of the basal ganglia. The Go cells disinhibit the thalamus via GPi, thereby facilitating the execution of an action represented in cortex. The NoGo cells have an opposing effect by increasing inhibition of the thalamus, suppressing actions from getting executed. Dopamine from the SNc projects to the dorsal striatum, causing excitation of Go cells via D1 receptors, and inhibition of NoGo via D2 receptors. GPi: internal segment of globus pallidus; GPe: external segment of globus pallidus; SNc: substantia nigra pars compacta; SNr: substantia nigra pars reticulata. **b)** The Frank (in press) neural network model of this circuit (squares represent units, with height and color reflecting neural activity; yellow = most active, red = less active, grey = not active). The Premotor Cortex selects an Output response via direct projections from the sensory Input, and is modulated by the BG projections from Thalamus. Go units are in the left half of the Striatum layer; NoGo in the right half, with separate columns for the two responses [R1 (left button), R2 (right button)]. In the case shown, striatum Go is stronger than NoGo for R1, inhibiting GPi, disinhibiting Thalamus, and facilitating execution of the response in cortex. A tonic level of dopamine is shown in SNc; a burst or dip ensues in a subsequent error feedback phase (not shown), causing corresponding changes in Go/NoGo unit activations, which drive learning. **c)** Predictions from the model for the probabilistic selection task, showing Go-NoGo associations for stimulus A and NoGo-Go associations for stimulus B. Error bars reflect standard error across 25 runs of the model with random initial weights.

1a

Probabilistic Selection

A (80%) B (20%)

C (70%) D (30%)

E (60%) F (40%)

Transitive Inference

A+ B-

B+ C-

C+ D-

D+ E-

1b

Probabilistic Selection

Test Performance

Transitive Inference
Test Performance

1d

Across Both Tasks

Z-scores

2a

Probabilistic Selection

BG Model Go/NoGo Associations