

Models of probabilistic category learning in Parkinson's disease: Strategy use and the effects of L-dopa

Maarten Speekenbrink^{a,*} David A. Lagnado^a Leonora Wilkinson^b
Marjan Jahanshahi^b David R. Shanks^a

^a*Cognitive, Perceptual and Brain Sciences, University College London*

^b*Sobell Department of Motor Neuroscience and Movement Disorders, Institute of Neurology,
University College London*

Abstract

Probabilistic category learning (PCL) has become an increasingly popular paradigm to study the brain bases of learning and memory. It has been argued that PCL relies on procedural habit learning, which is impaired in Parkinson's Disease (PD). However, as PD patients were typically tested under medication, it is possible that L-dopa caused impaired performance in PCL. We present formal models of rule-based strategy switching in PCL to re-analyse the data from Jahanshahi et al. (2009) comparing PD patients on and off medication (within subjects) to matched controls. Our analysis shows that PD patients followed a similar strategy switch process as controls when off medication, but not when on medication. On medication, PD patients mainly followed a random guessing strategy, with only few switching to the better Single Cue strategies. PD patients on medication and controls made more use of the optimal Multi-Cue strategy. In addition, while controls and PD patients off medication only switched to strategies which did not decrease performance, strategy switches of PD patients on medication were not always directed as such. Finally, results indicated that PD patients on medication responded according to a probability matching strategy indicative of associative learning, while the behaviour of PD patients off medication and controls was consistent with a rule-based hypothesis testing procedure.

Key words: Parkinson's disease, Levodopa, probabilistic category learning, strategy analysis, Levodopa

* Correspondence address: M. Speekenbrink, Cognitive Perceptual and Brain Sciences, University College London, 26 Bedford Way, London WC1H 0AP, United Kingdom.

Email addresses: m.speekenbrink@ucl.ac.uk (Maarten Speekenbrink), d.lagnado@ucl.ac.uk (David A. Lagnado), l.wilkinson@ion.ucl.ac.uk (Leonora Wilkinson), m.jahanshahi@ion.ucl.ac.uk (Marjan Jahanshahi), d.shanks@ucl.ac.uk (David R. Shanks).

URL: <http://www.ucl.ac.uk/psychlangsci/staff/cpb-staff/m.speekenbrink> (Maarten)

1. Introduction

Parkinson’s disease (PD) is well-known for its characteristic motor symptoms such as rigidity, tremor and akinesia (lack of movement). In addition to these, PD has been associated with a number of cognitive deficits. Dementia is prevalent in PD patients (Aarsland et al., 2005), but there is a range of less severe cognitive symptoms of PD, related to working memory (Owen et al., 1997), attention (Filoteo and Maddox, 1999), set-shifting (Cools et al., 2001b), and procedural learning (Saint-Cyr et al., 1988). PD patients have also been shown to be impaired in category learning (Ashby et al., 2003; Filoteo et al., 2005, 2007; Knowlton et al., 1996; Maddox et al., 2005; Shohamy et al., 2004b). Given the nature of the disease, it may come as no surprise that PD patients show certain learning deficits. PD involves the loss of dopamine producing cells in the substantia nigra pars compacta, resulting in dopamine depletion in the dorsal striatum, extending to other areas such as the ventral striatum and prefrontal cortex as the disease progresses (Cools, 2006). Dopamine is thought to play a crucial role in the flexible control of behaviour in response to environmental demands (Cools, 2006; Nieoullon, 2002). Evidence suggests that phasic dopamine bursts and dips in the striatum code “reward prediction errors” (O’Doherty et al., 2004; Schultz, 2002; Schultz et al., 1997), associated with the presence of an unexpected reward and absence of an expected reward respectively. Prediction errors are crucial to many learning theories and the disruption of these signals in PD is thus likely to impair learning.

In this article, we restrict our attention to Probabilistic Category Learning (PCL), involving tasks in which stimulus features are imperfect predictors of category membership. A popular PCL task is the Weather Prediction Task (WPT, Knowlton et al., 1994), which asks participants to predict the state of the weather (Rainy, or Fine) based on combinations of four “tarot” cards which are related to the weather with different probabilities. Research with the WPT has provided mixed results. Some studies showed that, compared to healthy controls, PD patients were impaired at early learning (Knowlton et al., 1996; Witt et al., 2002), but to a lesser extent (Knowlton et al., 1996) or not (Witt et al., 2002) in later learning. Other studies found impaired performance throughout the task (Shohamy et al., 2004b,a; Wilkinson et al., 2008). Some studies found no impairment in performance at all (Moody et al., 2004; Price, 2005).

One explanation of these discrepant findings is that there are different ways in which PCL tasks such as the WPT can be solved, and that these different learning strategies implicate distinct neural circuits. One impetus for the claim of such dissociable learning systems was the finding that amnesic patients showed no (early) impairment in the WPT, whilst unable to recall certain aspects of the testing episode (Eldridge et al., 2002; Knowlton et al., 1994, 1996). Impaired WPT performance was not accompanied by this lack of declarative memory in PD patients (Knowlton et al., 1996; Sage et al., 2003; Witt et al., 2002). The double dissociation between learning performance and explicit recall has been taken as evidence for the existence of multiple memory systems (Ashby et al., 1998; Ashby and Maddox, 2005; Gabrieli, 1998; Knowlton et al., 1994, 1996; Poldrack and Rodriguez, 2004; Shohamy et al., 2008). According to this view, PCL involves an implicit habit learning process which depends on the dorsal striatum, an area unaffected

Speekenbrink).

by amnesia. Declarative memory on the other hand involves the medial temporal lobes, an area which is impaired in amnesic, but not PD patients. Knowlton et al. (1996) suggested that PD patients who did learn in the WPT relied on declarative learning strategies, which is supported by the finding that PD patients showed normal learning in a paired-associate version of the WPT, emphasizing declarative learning by memorization (Shohamy et al., 2004a). This suggests that PD patients have particular problems in learning from feedback to their responses, which can be related to disrupted reward prediction errors. Neuroimaging studies with healthy individuals have shown increased activity in the striatum and decreased activity in the medial temporal lobes during the normal (feedback based) WPT (Poldrack et al., 2001; Witt et al., 2002). PD patients on the other hand show increased activation in the medial temporal lobes throughout the task, as well as lower activity in the dorsal striatum than controls (Witt et al., 2002). The emerging picture is thus that, due to a disrupted striatal habit learning system, PD patients rely on a qualitatively different learning process, one that is explicit (declarative) rather than implicit (procedural).

Upon closer scrutiny, there is reason to doubt the implicit, habitual nature of PCL. Both healthy (Lagnado et al., 2006) and amnesic (Speekenbrink et al., 2008) participants have shown explicit insight into the cue-outcome contingencies, thus indicating access to representations of the task environment which should be absent in purely habitual learning. Indeed, there is now considerable evidence that PCL relies at least partially (Meeter et al., 2006), if not wholly (Lagnado et al., 2006; Newell et al., 2007; Price, 2009), on explicit processes. Furthermore, a recent study failed to replicate the results of Shohamy et al. (2004a). Instead of selective impairment for the feedback-based, but normal learning on the paired associate version of the WPT, we found PD patients were impaired on both versions relative to controls (Wilkinson et al., 2008). With the above in mind, it is important to note that previous studies have almost exclusively tested PD patients whilst on medication. The current medication levodopa (L-dopa) is a precursor for dopamine and raises the dopamine levels in affected areas. Whilst generally effective in reducing motor symptoms, the effects of L-dopa on cognitive functioning are less clear, sometimes improving and sometimes deteriorating cognitive ability (Cools et al., 2001a; Swinson et al., 2000; Gotham et al., 1988). Evidence suggests that the relation between dopamine and performance follows an “inverted U-shaped” function (Cools, 2006). Increasing dopamine levels can either enhance or lower performance, depending on baseline dopamine levels, implying a possible L-dopa “overdose” effect on relatively spared areas such as the ventral striatum and/or the prefrontal cortex (Cools et al., 2001a; Cools, 2006; Frank, 2005; Gotham et al., 1988). This raises the question whether the impairment in PCL displayed by PD patients is the result of impaired habit learning due to dorsal striatal dysfunction, or of an L-dopa overdose effect on relatively intact areas. As the ventral striatum and prefrontal cortex are thought to support more intentional and representational forms of learning, impaired PCL learning by PD patients on medication may be due to impaired explicit rather than implicit learning processes.

In this article, we will apply a formal modelling framework to assess the nature of PCL in PD and how it is affected by L-dopa. Based on work by Gluck et al. (2002), we formulate a model of rule-based learning consistent with explicit learning strategies such as hypothesis testing. We contrast it to a model which is more closely related to associative or habitual learning. We apply the models to data from Jahanshahi et al. (2009), who conducted an experiment in which PD patients performed the WPT both on and

off medication (L-dopa). This data allows us to test whether PD patients approach PCL tasks in a different way than healthy controls and whether their learning strategy is affected by L-dopa. If performance in the WPT depends primarily on a habit learning system supported by the dorsal striatum (e.g., Knowlton et al., 1994, 1996), we should expect the associative model to fit control participants better than the rule-based model. We should expect the opposite pattern for PD patients, who must overcome impaired habit learning by relying on declarative strategies. If, on the other hand, performance depends on more explicit processes, we should expect the rule-based model to fit both controls and PD patients better than the associative model. However, if an L-dopa overdose effects impairs these explicit strategies, we might expect the rule-based model to fit poorly to PD patients on medication.

In the following, we will first describe the Weather Prediction Task, followed by a brief description of the the study conducted by Jahanshahi et al. (2009). We will then introduce our strategy switch model and describe the different versions that we fitted to the data. In the results section, we first discuss how these models described behaviour in the PD and control group separately, and then investigate group differences in learning strategies.

2. The Weather Prediction Task

In the Weather Prediction Task, participants are asked to predict the state of the weather Y (1 = Fine, 0 = Rainy), on the basis of four “tarot cards” x_j (cards with geometrical patterns), which are either presented ($x_j = 1$) or not ($x_j = 0$). Each card is associated with the weather with a different probability. Jahanshahi et al. (2009) used the version introduced by Gluck et al. (2002), in which Card 1 (card 4) is strongly ($p = .8$) associated with Fine (Rainy) weather, and card 2 (card 3) is weakly ($p = .6$) associated with Fine (Rainy) weather. On each learning trial, participants are presented with a combination of one, two, or three cards. The probability of the card combinations is given in Table 1. Participants are then asked to predict the state of the weather, after which they receive immediate outcome feedback on the actual state of the weather. In addition, Jahanshahi et al. (2009) informed participants whether the response was correct or incorrect, by a “thumbs up” and “thumbs down” sign respectively. The state of the weather depends on the 14 possible cue patterns $\mathbf{x} = (x_1, \dots, x_4)$ according to the conditional probabilities in Table 1. The objective is to draw upon these different predictive patterns to maximize the number of correct predictions.

3. The Jahanshahi et al. (2009) study

Jahanshahi et al. (2009) conducted a study to investigate the effects of L-dopa on PCL. They employed a within subjects design, in which PD patients were tested with the WPT both on and off L-dopa, and matched controls also performed the task twice. As a detailed description of the study can be found in Jahanshahi et al. (2009), we only provide a brief description of the relevant details here.

Table 1

Structure of the Weather Prediction task and strategy definitions ($v_s(\mathbf{x})$) for the Constant Error and Generalized Matching strategy switch model.

	Pattern													
	A	B	C	D	E	F	G	H	I	J	K	L	M	N
	\mathbf{x}													
	0001	0010	0011	0100	0101	0110	0111	1000	1001	1010	1011	1100	1101	1110
$P(\text{Pattern})$	0.095	0.045	0.130	0.045	0.060	0.030	0.095	0.095	0.030	0.060	0.045	0.130	0.045	0.095
$P(\text{Fine} \text{Pattern})$	0.895	0.778	0.923	0.222	0.833	0.500	0.895	0.105	0.500	0.167	0.556	0.077	0.444	0.105
Constant Error (CE)														
Random	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Singleton	1	1	0	-1	0	0	0	-1	0	0	0	0	0	0
Single 1	1	1	1	1	1	1	1	-1	-1	-1	-1	-1	-1	-1
Single 2	1	1	1	-1	-1	-1	-1	1	1	1	1	-1	-1	-1
Single 3	-1	1	1	-1	-1	1	1	-1	-1	1	1	-1	-1	1
Single 4	1	-1	1	-1	1	-1	1	-1	1	-1	1	-1	1	-1
Multi-cue	1	1	1	-1	1	0	1	-1	0	-1	1	-1	-1	-1
Generalized Matching (GM)														
Random	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Singleton	2.143	1.254	0	-1.254	0	0	0	-2.143	0	0	0	0	0	0
Single 1	1.386	-1.386	1.386	-1.386	1.386	-1.386	1.386	-1.386	1.386	-1.386	1.386	-1.386	1.386	-1.386
Single 2	-0.405	0.405	0.405	-0.405	-0.405	0.405	0.405	-0.405	-0.405	0.405	0.405	-0.405	-0.405	0.405
Single 3	0.405	0.405	0.405	-0.405	-0.405	-0.405	-0.405	0.405	0.405	0.405	0.405	-0.405	-0.405	-0.405
Single 4	1.386	1.386	1.386	1.386	1.386	1.386	1.386	-1.386	-1.386	-1.386	-1.386	-1.386	-1.386	-1.386
Multi-cue	2.143	1.254	2.484	-1.254	1.607	0	2.143	-2.143	0	-1.607	0.225	-2.484	-0.225	-2.143

3.1. Participants

The PD patient group consisted of eleven individuals (8 male, 3 female), aged between 53 and 73 ($M = 63.5$, $SD = 6.2$), with a diagnosis of idiopathic Parkinson’s disease. Patients were recruited from the movement disorders clinic at the National Hospital for Neurology and Neurosurgery and their disease duration ranged from 3 to 37 years ($M = 13.2$, $SD = 10.7$). The majority of patients were in the mild to moderate stages of the disease, while the disease was more severe for a small number of patients. When patients were assessed off medication, the mean Hoehn-Yahr score (Hoehn and Yahr, 1967) was 2.8 ($SD = 0.6$) and when assessed on medication, the mean Hoehn-Yahr score was 2.2

($SD = 0.8$). All patients were non-demented and did not suffer from clinical depression. In addition, all patients in the study were treated with L-dopa, stable on their medication doses and responding well to the medication. The control group consisted of thirteen healthy volunteers (5 male, 8 female) aged between 44 and 69 ($M = 60.00$, $SD = 9.7$). Controls were screened and none had any neurological disorder or history of psychiatric illness.

3.2. Procedure

PD patients were tested on two consecutive days. Half of the PD patients were tested on medication first and the remainder were tested off medication first. When patients were assessed on medication they were tested while they continued to take their usual levodopa medication, on average 1.69 ($SD = 0.94$) hours after taking their last dose of medicine. When assessed off medication, patients were tested after overnight withdrawal of medication, on average 13.79 ($SD = 2.2$) hours after taking their last dose of medication. Controls completed both tasks on the same day. Because each participant did the WPT under two different conditions (on/off medication or time 1/time 2), a parallel version of the usual WPT was constructed in which the outcomes were labelled Hot and Cold. In addition, four different sets of tarot cards were used. Each version of the WPT consisted of 200 trials. The order of the two versions was randomized.

3.3. Learning performance

To assess overall differences in learning performance, we re-analysed the results using the average probability of correct responses, rather than the proportion of optimal responses, as a measure of performance (see Speekenbrink et al., 2008, for a specification and justification of this measure). A 2 (group) by 2 (occasion) by 4 (block) ANOVA, with repeated measures on the last two factors, showed a significant main effect of block, $F(3, 66) = 4.54$, $MSE = 0.004$, $p = .006$, and a marginally significant main effect of occasion, $F(1, 22) = 4.06$, $MSE = 0.01$, $p = .056$. In general, performance increased during the task, block 1: $M = 0.62$, $SD = 0.32$; block 2: $M = 0.64$, $SD = 0.31$; block 3: $M = 0.66$, $SD = 0.30$; block 4: $M = 0.67$, $SD = 0.30$. Furthermore, performance increased from the first ($M = 0.63$, $SD = 0.31$) to the second testing occasion ($M = 0.66$, $SD = 0.30$). There were no other significant effects. Importantly, although PD patients ($M = 0.63$, $SD = 0.32$) performed slightly less well than controls ($M = 0.67$, $SD = 0.30$), this difference was not significant, $F(1, 22) = 2.42$, $MSE = 0.028$, $p = .134$. To investigate the effect of medication on performance, we conducted a 2 (order: on medication first, off medication first) by 2 (medication: on, off) by 4 (block) ANOVA on the PD patient data, with repeated measures on the last two factors. This showed a significant main effect of medication, $F(1, 9) = 5.58$, $p = .042$, $MSE = 0.002$. PD patients performed better off ($M = 0.64$, $SD = 0.31$) than on medication ($M = 0.62$, $SD = 0.32$). Other effects were not significant. Taken together, these results indicate that PD patients were impaired when on medication, but not when off medication. This was confirmed by t-tests, which showed that performance of PD patients on medication differed from controls, $t(22) = 2.12$, $p < .05$, while the performance of PD patients off medication did not differ from controls, $t(22) = 0.87$, $p = 0.39$.

4. Strategy switch model

Our strategy switch model is based on the work of Gluck et al. (2002). On the basis of participants' answers in a post-task questionnaire, Gluck et al. distinguished three broad classes of response strategies in the WPT: a Singleton strategy, where participants respond optimally to those cue patterns consisting of a single card, but guess for combinations of multiple cards, Single Cue strategies, in which participants base their responses solely on the presence or absence of a single card, ignoring the other cues, and a Multi-Cue strategy, in which participants learn to respond optimally to all possible card patterns. Gluck et al. (2002) fitted the strategies to participants' responses by least squares and showed that they adequately fitted the large majority of their participants. This strategy analysis has proven popular in neuropsychological investigations of PCL (Fera et al., 2005; Shohamy et al., 2004b,a; Wilkinson et al., 2008). Previous results have shown that healthy participants typically use the Singleton strategy early on in the task, moving to the Single Cue and Multi-Cue strategies at later stages (Gluck et al., 2002). PD patients, on the other hand, appear to rely on simple strategies such as the Singleton and Single Cue strategies throughout the task (Shohamy et al., 2004b,a).

There are different views on what shifts in strategy use signify. One interpretation (e.g., Ashby and Maddox, 2005) is that strategy shifts indicate switches between qualitatively different learning processes, such as direct memorization (Singleton strategy), rule-learning (Single Cue strategies) and incremental associative learning (Multi-Cue strategy). Another interpretation (e.g., Meeter et al., 2008) is that they indicate the adoption of different hypothetical rules in a hypothesis testing procedure. Another interpretation (e.g., Shohamy et al., 2008) is that strategy switches reflect the progression of an incremental associative learning procedure, in which cue-outcome (or stimulus-response) associations slowly adapt to the task structure. Responses become more optimal as a result, which may lead to an apparent shift from suboptimal Single Cue strategies to the optimal Multi-Cue strategy. The original strategy analysis offers little guidance to discriminate between these alternative interpretations. Rules are typically deterministic, while the task representations resulting from associative learning will reflect the probabilistic nature of the task. Assuming that responses are made in accordance with this representation, we may expect responses from associative learning to reflect the cue-outcome probabilities. Based on this distinction, we formulate two variants of the strategies: one in which responses are made according to deterministic rules, allowing for a fixed level of "implementation error", and one in which responses are made by matching to the conditional probabilities of the outcome.

Shifts in strategy use are usually identified by fitting strategies to separate blocks of trials. More recently, Meeter et al. (2006) improved upon this method by specifically focussing on strategy switches and attempting to identify the switch trials more precisely. We build on this idea here, but rather than assigning maximum likely strategies to small blocks of trials (Meeter et al., 2006, 2008), we formalize the strategy switch model as a hidden Markov model. This offers a more principled approach with a number of advantages. For instance, we can obtain maximum likelihood estimates of the probabilities of switching between strategies. These can be used to test for group differences in the strategy switching process. By restricting transition probabilities between certain strategies, we can also assess the rationality of the strategy switch process. Furthermore, we can use

the model to determine, for each participant, a maximum a posteriori strategy sequence. This offers a statistically optimal way of assigning strategy sequences to participants.

4.1. General structure

Meeter et al. (2006) formulated 11 strategies on the basis that each (1) should result in above chance performance, and (2) should be easily verbalisable. As the difference between some of the proposed strategies is rather subtle, we restrict our attention to a subset of 7, which contains the original strategies identified by Gluck et al. (2002). In words, these strategies can be stated as:

- (i) *Random*: Guess (on each trial, predict “Fine” with probability .5, and “Rainy” otherwise. Note that this strategy does not result in above chance performance. It is included to allow more valid assignment of the remaining strategies).
- (ii) *Singleton*: predict “Fine” when either only cue 1 or cue 2 is present, predict “Rainy” when either only cue 3 or cue 4 is present, and guess otherwise.
- (iii) *Single Cue 1*: Predict “Fine” when cue 1 is present, and “Rainy” otherwise.
- (iv) *Single Cue 2*: Predict “Fine” when cue 2 is present, and “Rainy” otherwise.
- (v) *Single Cue 3*: Predict “Rainy” when cue 3 is present, and “Fine” otherwise.
- (vi) *Single Cue 4*: Predict “Rainy” when cue 4 is present, and “Fine” otherwise.
- (vii) *Multi-cue*: Predict “Rainy” when only cue 4 is present, . . . , predict “Rainy”, when cue 3 and 4 are present, . . . , predict “Fine” when cue 1, 2, and 3 are present.

As the name suggests, according to the strategy switch model, participants can change their strategy on a trial by trial basis. We formalize the strategy switch model as a hidden Markov model (e.g., Rabiner, 1989; Wickens, 1982). Hidden Markov models describe the inter-dependencies in a series of observations through a latent Markovian state process, which in the present context corresponds to the strategies used. Hidden Markov models are characterized by two assumptions: (1) responses R_t are conditionally independent given states S_t ,

$$P(R_1, R_2, \dots, R_T | S_1, S_2, \dots, S_T) = \prod_{t=1}^T P(R_t | S_t), \quad (1)$$

and (2) states S_t depend only on the previous state S_{t-1} (Markov property),

$$P(S_1, \dots, S_T) = P(S_1) \prod_{t=2}^T P(S_t | S_{t-1}). \quad (2)$$

In addition, we will assume the process is stationary, or time homogeneous, so that $P(R_t = r | S_t = s) = P(R_{t-1} = r | S_{t-1} = s)$ and $P(S_{t+1} = s | S_t = q) = P(S_t = s | S_{t-1} = q)$ for all q, r, s, t . To complete the model, we need to specify the conditional distributions $P(S_{t+1} | S_t)$ and $P(R_t | S_t)$. For the latter, we use a Bernoulli distribution

$$P(R_t = r | S_t = s) = p_s(\mathbf{x}_t)^r [1 - p_s(\mathbf{x}_t)]^{1-r}, \quad (3)$$

with $r \in \{0, 1\}$. The cue vectors \mathbf{x}_t are treated as fixed and, for ease of notation, we will usually not explicitly conditionalize on them.

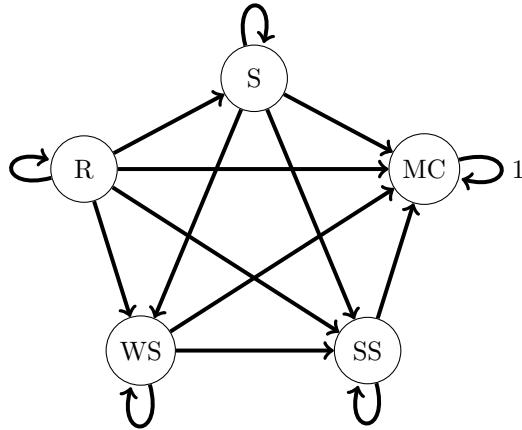


Figure 1. State transitions in the Directed Transition (DT) strategy switch models. Non-zero transition probabilities are depicted as arrows. R = Random, S = Singleton, WS = Weak Single Cue, SS = Strong Single Cue, MC = Multi-Cue strategy.

4.2. Undirected and directed transitions

The (stationary) transition probabilities $a_{ij} = P(S_{t+1} = j | S_t = i)$ are collected in the transition matrix $\mathbf{A} = \{a_{ij}\}, i, j = 1, \dots, M$. In one version of the model, we allow transitions between all M states. For example, individuals can switch from the optimal Multi-Cue strategy to the rather poor Random strategy. We will refer to these as the Undirected Transition (UT) models. We contrast these to the theoretically interesting Directed Transition (DT) models, in which participants can only switch to strategies which are at least as good as the previous strategy in terms of expected performance. As such, the DT version implements a kind of “rationality constraint”. In the DT version, we also treat the Single Cue 1 and Single Cue 4 strategies (strong cues), and the Single Cue 2 and Single Cue 3 strategies (weak cues) as similar. In particular, we assume that the probability of switching to and from the two strong (weak) Single Cue strategies are identical to each other. A graphical representation of the resulting state transition model is given in Figure 1. In the DT models, the transition matrix \mathbf{A} has 12 free parameters (compared to the 42 free parameters of the UT models). Note that the Multi-Cue strategy is an absorbing state (e.g., Wickens, 1982) – as it results in the best performance, participants cannot switch once it has been adopted.

4.3. Constant Error model

Meeter et al. (2006) assumed that the strategies were implemented probabilistically, such that on each trial, there is a fixed probability p that the response is in accordance with the strategy, and a probability $1 - p$ that the response is the opposite of that prescribed by the strategy. This is consistent with rule-following behaviour with a fixed level of “implementation error”. In their study, they used values of $p = .9$, $p = .95$ and $p = .99$. Rather than fixing the probability to such pre-specified values, we let p be a freely estimable parameter in the range $p \in [.8; 1]$. The lower bound of this range was

chosen to increase the discrimination between the strategies. Obviously, the strategies are indistinguishable if $p = .5$ (and $p < .5$ would mean that individuals can apply a strategy whilst being more likely to respond in disagreement than in agreement with the strategy). The higher the value of p , the more the strategies differ in their predicted responses. While the exact value of the lower bound ($p \geq .8$) is somewhat arbitrary, we felt that lower values of p would allow the randomness in responses to exceed plausible levels for rule-following behaviour.

We formulate the response part of the model as a generalized linear model (Dobson, 2002; McCullagh and Nelder, 1983), using Equation 3 with

$$p_s(\mathbf{x}) = h(\lambda v_s(\mathbf{x})), \quad (4)$$

where

$$h(z) = \frac{\exp(z)}{1 + \exp(z)} \quad (5)$$

is the inverse of the logit link function

$$g(p) = \ln(p/(1 - p)), \quad 0 < p < 1. \quad (6)$$

Equations 4 and 5 specify a logistic regression model with a single predictor $v_s(\mathbf{x})$. The cue-dependent values of the predictor $v_s(\mathbf{x})$ are given in Table 1. As can be seen there, $v_s(\mathbf{x})$ can only take the value -1, 0, or 1 in the Constant Error model. With these values, the regression coefficient λ is related to the probability of a strategy consistent response as

$$p = h(\lambda). \quad (7)$$

The lower bound $p \geq 0.8$ corresponds to a lower bound of $\lambda \geq 1.386$ (there is no upper bound for λ ; as λ increases to infinity, p approaches 1). While the formulation as a logistic regression model is not strictly necessary at this point, it is useful for the alternative version of the strategy switch model we present next.

4.4. Generalized Matching model

The Constant Error model assumes that participants follow a strategy with a certain fixed probability p , which is identical for each strategy. We now propose an alternative model in which $p_s(\mathbf{x})$ depends on the predictive validity (or diagnosticity) of the different cue patterns. This model is related to the ‘‘multi-match’’ strategy identified by Lagnado et al. (2006). The multi-match strategy is like the optimal Multi-Cue strategy, but responses are made by probability matching to the conditional probabilities $P(Y|\mathbf{x})$. The original Multi-Cue strategy of Gluck et al. (2002) predicts maximising responses to all cue patterns, regardless of the diagnosticity of the pattern. Similarly, according to the Multi-Cue strategy in the CE model, the probability of a maximising response is identical for each cue pattern. For example, the probability of a maximising response is identical for responses to cue pattern A and cue pattern K. As can be seen in Table 1, pattern A is relatively predictive of Fine weather ($P(\text{Fine}|A) = 0.895$), while the predictive validity of pattern K is relatively low ($P(\text{Fine}|K) = 0.556$). As such, it is likely that inferring the optimal response is easier for pattern A than for pattern K. This difference can be incorporated into the model by letting the probability of responses vary according to the diagnosticity of the cues. We extend this idea to all strategies, so that, for example,

responses are made in accordance with $P(Y|x_1)$ in the Single Cue 1 strategy. Moreover, we use a response scaling parameter $\lambda \geq 1$ which allows responses to vary between probability matching and maximising. We use the exponential version of Luce’s (Luce, 1959) choice rule, which is commonly used in categorization models (e.g., Ashby and Maddox, 1993). For example, the generalized matching version of the Multi-Cue strategy can be stated as

$$P(R = r|S = \text{Multi-Cue}) = \frac{P(Y = r|\mathbf{x})^\lambda}{P(Y = r|\mathbf{x})^\lambda + P(Y \neq r|\mathbf{x})^\lambda}. \quad (8)$$

By changing the conditional term in Equation 8, we can derive the response rule for the other strategies (e.g., replacing \mathbf{x} for the Multi-Cue strategy by x_1 for the Single Cue 1 strategy). In the GM version of the strategy switch model, responses are made in accordance with the (conditional) probability of the criterion Y , and strategies determine which cues function in the conditional. In this way, a strategy might reflect the attention given to different cues (assuming attention is an all-or-none process). This generalized matching version of the strategy switch model is easily implemented in the current framework by using

$$v_s(\mathbf{x}) = g(P(Y = 1|\cdot)) \quad (9)$$

in Equation 4, where g is defined as in Equation 6 and the dot replaced by the relevant conditional term. The values of $v_s(\mathbf{x})$ for the GM model are given in Table 1.

The assumed relation between response and criterion probabilities in Equation 8 and 9 is a special case of the generalized matching law (Baum, 1974, 1979). The generalized matching law has provided a successful description of responses in a wide range of operant conditioning studies, both for human (see e.g. Kollins et al., 1997, for an overview) and non-human (Baum, 1979) animals. In previous work (Speekenbrink et al., 2008), we formulated an incremental, associative learning model – closely related, but not identical to the Rescorla-Wagner (Rescorla and Wagner, 1972) model – which predicts responses in accordance with this generalized matching law. In this respect, the Multi-Cue strategy in the GM version corresponds to the outcome of an incremental associative learning model. Likewise, the single cue strategies correspond to the outcome of incremental learning when the learner focusses attention to a single cue. Switches between the strategies can thus correspond to rapid shifts in attention (cf. Kruschke and Johansen, 1999).

4.5. Estimation

The strategy switch model contains a number of freely estimable parameters. These are the initial state probabilities $P(S_1 = j)$, the transition probabilities a_{ij} , and the response consistency parameter λ . For simplicity, and theoretical reasons, we fixed the initial state probabilities so that $P(S_1 = \text{Random}) = 1$. The remaining parameters are estimated by the Expectation-Maximisation (EM, Dempster et al., 1977) algorithm. Details of the estimation procedure are provided in the Appendix. There, we also give details of the procedure to determine the Maximum A Posterior (MAP) state sequences, the most probable strategy sequences given the data and model parameters.

4.6. Inference

The strategy switch model can be used to investigate group differences by comparing versions with different equality restrictions on the model parameters. For instance, we can restrict the transition matrix \mathbf{A} to be identical for PD patients on and off medication, or for PD patients off medication and controls. Versions with such equality restrictions are nested under versions without them, and as the model parameters are estimated by maximum likelihood, we can use likelihood ratio tests to assess whether the restrictions significantly impair model fit. The likelihood ratio test statistic is computed as

$$LRT = 2[\ln(L_1) - \ln(L_2)], \quad (10)$$

where $\ln(\cdot)$ is the natural logarithm, L_j the model likelihood and k_j the number of freely estimated parameters of model $j = 1(\text{general}), 2(\text{nested})$. The LRT statistic is asymptotically Chi-square distributed, with $k_1 - k_2$ degrees of freedom (e.g., Dobson, 2002).

To compare non-nested models (such as the UT and DT models, where the latter fixes certain parameters of the former at 0, which is on the bound of the parameter space), we can use model selection criteria such as the AIC (?) and BIC (Schwarz, 1978) as a means to infer the tenability of parameter restrictions.

4.7. Relation to other models

A large number of formal learning models have been proposed to describe probabilistic category learning (see e.g. Ashby and Maddox, 2005, for a partial overview). The objective of these models is to describe how participants learn from the information they have been given. Transitions from one model state to the next are an essentially deterministic function of outcome feedback. While participants' responses are used to adjust model parameters such as learning rate and response consistency, these adjustments are often a matter of degree. For example, in the well known Rescorla-Wagner (Rescorla and Wagner, 1972) model, the change in associative strength between a cue and outcome is proportional to the difference between observed and predicted outcome. While the learning rate determines the scale of the difference (i.e., by how much associative strength changes) it is the outcome that determines the direction of the difference (i.e., whether associative strength increases or decreases). In this sense, the outcome feedback constrains the path the model takes. A different approach is to estimate model states directly from participants' responses, as in rolling regression (Kelley and Friedman, 2002; Lagnado et al., 2006) and dynamic lens model analysis (Speekenbrink et al., 2008; Speekenbrink and Shanks, 2008). These models make a similar assumption regarding the relation between responses and cues as the Rescorla-Wagner model (e.g., responses are an additive function of cue values), but estimate an associative-weight-like quantity, called cue utilization, without reference to the outcome feedback. As such, cue utilization can increase or decrease freely on each trial, allowing for idiosyncratic and sometimes quite irrational patterns. A main objective of dynamic lens model analysis is to compare these relatively unconstrained estimates of model states to the relatively constrained estimates of formal learning models in order to assess the latter's validity.

The strategy switch model lies between these two approaches. The model proposes a

number of a priori determined strategies, but makes no assumption how transitions between strategies are related to outcome feedback. Meeter et al. (2008) compared rolling regression to their version of the strategy switch model and concluded that both were about equal in their ability to predict participants' responses. In comparison to their strategy switch model, our version is both more and less constrained. As we include a freely estimable response consistency parameter, there is more freedom in fitting strategies to participants' responses. However, in our version, the strategy switch process depends on the transition matrix. Once estimated, this places constraints on strategy switches which are absent in the model of Meeter et al. (2008). Moreover, as implemented in the Directed Transition model, the transition probabilities can be fixed to place a priori restrictions on strategy switches, constraining strategy switches even more.

In comparison to other models of probabilistic category learning, our strategy switch model has a relatively large number of free parameters. However, apart from the response consistency parameter, they are all estimates of the transition probabilities between strategies and have a direct and clear interpretation. While we could have restricted the number of parameters by setting the transition probabilities equal for all strategies, we were particularly interested in group differences in strategy switches. By placing too much constraint on the transition matrix, the model may miss these differences.

5. Results

We fitted several versions of the strategy switch model, each with different restrictions on the parameters. We first fitted the different models to the PD and control group separately. For each of the four models (CE/UT, CE/DT, GM/UT and GM/DT), we constructed a series of nested versions by imposing different equality constraints on the model parameters. The transition matrix could vary over testing occasions (denoted as \mathbf{A}_o), or not (denoted as \mathbf{A}). In addition, response consistency could vary over participants and occasions (denoted as $\lambda_{i,o}$), vary over participants but not occasions (λ_i), vary over occasions but not participants (λ_o), or not vary at all (λ). The resulting fit measures for the models are given in Table 2.

5.1. PD patients

For the PD patients, the overall best fitting model according to the AIC was the GM/UT model with an identical transition matrix (\mathbf{A}) but different individual consistency parameter in the on and off medication conditions ($\lambda_{i,o}$). The BIC selected as overall best fitting the simpler GM/DT model with identical transition probabilities (\mathbf{A}) and consistency parameters (λ) for the on and off medication conditions. Disagreement between the selection criteria is not uncommon. As sample size increases, the BIC puts a higher penalty on additional parameters than the AIC. In fact, as can be seen in Table 2, the BIC consistently selected the model with the least number of parameters, which may be due to the relatively large sample size (there are 4400 observations in the PD conditions).

For a more overall comparison of the different models, we used Wilcoxon tests to com-

Table 2
Fit measures for the Strategy Switch models. Best fitting models in bold.

parameters	# par.	Constant error			Generalized Matching		
		log L	AIC	BIC	log L	AIC	BIC
PD patients							
Undirected Transition (UT)							
A λ	43	-2401	4889	5164	-2397	4879	5154
λ_i	53	-2388	4882	5221	-2389	4884	5222
λ_o	44	-2395	4878	5160	-2401	4889	5170
$\lambda_{i,o}$	64	-2380	4887	5296	-2372	4873	5282
A_o λ	85	-2375	4920	5463	-2366	4902	5445
λ_i	95	-2364	4918	5525	-2365	4920	5527
λ_o	86	-2375	4922	5471	-2376	4923	5473
$\lambda_{i,o}$	106	-2352	4915	5593	-2343	4898	5575
Directed Transition (DT)							
A λ	13	-2449	4925	5008	-2436	4897	4980
λ_i	23	-2432	4909	5056	-2427	4901	5048
λ_o	14	-2446	4921	5010	-2435	4899	4988
$\lambda_{i,o}$	34	-2424	4916	5133	-2420	4909	5126
A_o λ	25	-2446	4941	5101	-2430	4910	5069
λ_i	35	-2428	4926	5150	-2422	4913	5137
λ_o	26	-2442	4937	5103	-2429	4911	5077
$\lambda_{i,o}$	46	-2421	4933	5227	-2416	4925	5219
Controls							
Undirected Transition (UT)							
A λ	43	-2516	5118	5399	-2543	5173	5455
λ_i	55	-2506	5123	5483	-2528	5166	5527
λ_o	44	-2513	5114	5402	-2533	5155	5443
$\lambda_{i,o}$	68	-2479	5094	5539	-2510	5157	5602
A_o λ	85	-2477	5125	5682	-2523	5216	5773
λ_i	97	-2464	5123	5759	-2507	5208	5844
λ_o	86	-2477	5125	5689	-2523	5218	5782
$\lambda_{i,o}$	110	-2445	5111	5832	-2500	5219	5941
Directed Transition (DT)							
A λ	13	-2548	5122	5207	-2555	5137	5222
λ_i	25	-2531	5112	5276	-2542	5135	5298
λ_o	14	-2545	5118	5210	-2553	5133	5225
$\lambda_{i,o}$	38	-2506	5089	5338	-2522	5119	5369
A_o λ	25	-2543	5136	5300	-2548	5145	5309
λ_i	37	-2529	5133	5375	2525	5123	5366
λ_o	26	-2537	5125	5296	-2546	5145	5315
$\lambda_{i,o}$	50	-2496	5091	5419	-2510	5119	5447

Table 3
 Model fits for the PD on and PD off medication conditions separately.

model	# par.	PD on			PD off		
		log L	AIC	BIC	log L	AIC	BIC
CE/UT	53	-1217	2541	2843	-1134	2375	2677
CE/DT	23	-1260	2566	2697	-1161	2367	2498
GM/UT	53	-1210	2526	2828	-1133	2372	2674
GM/DT	23	-1255	2555	2686	-1162	2369	2500

pare the AIC and BIC values for corresponding versions of each model¹. This confirmed that the GM family was preferred to the CE family, $Z = 2.95$, $p < .01$. The AIC showed a general preference for the UT family, $Z = 2.12$, $p < .05$, while the BIC values indicated superiority of the DT family, $Z = -3.52$, $p < .001$.

For each of the four models, we used likelihood ratio tests to identify group differences in model parameters. For the GM/UT model, the on and off medication conditions differed in state transition probabilities, $\chi^2(42) = 58.98$, $p < .05$. Besides a difference in transition probability, there was clear evidence of individual variability in response consistency, $\chi^2(20) = 65.65$, $p < .001$, and that individual response consistency varied with medication, $\chi^2(11) = 44.12$, $p < .001$. The likelihood ratio tests thus confirm that the response consistency varied between individuals and conditions, but in addition also showed that the transition probabilities varied between conditions. This latter result is due to the greater power of the likelihood ratio test when comparing nested models.

The analyses above looked for an overall best fitting model for the PD patients as a group, and then assessed differences in parameters as a result of medication. This showed an overall preference for the GM/UT model, and that the PD on and off medication conditions differed in strategy transition probabilities and response consistency. To investigate whether medication was related to use of the GM or CE strategies, we looked at the fit measures for the \mathbf{A}_o and $\lambda_{i,o}$ versions of the four models separately for each condition. The results are given in Table 3. This analysis confirmed that the GM model fitted best for the PD patients on medication, with the AIC selecting the GM/UT model, and the BIC selecting the GM/DT model. For the PD patients off medication however, both the AIC and BIC selected the CE/DT model. The overall preference for the GM model noted earlier appears to be due to the patients on medication and for these participants, the superiority of the GM version is clearly marked.

To assess the particular strategies used, we computed the Maximum A Posteriori (MAP) strategy sequence for each PD patient, using the GM/UT model when on medication, and the CE/DT model when off medication. The results are depicted in Figure 2, where we classified the Single Cue strategies as “Strong Single Cue” strategies (Single

¹ We compared pairs of models with the same restrictions on \mathbf{A} and λ . For example, we compared the CE/DT model with \mathbf{A}_o and $\lambda_{i,o}$ to the corresponding GM/DT model with \mathbf{A}_o and $\lambda_{i,o}$. As the corresponding versions of the CE and GM models have the same number of parameters, differences in the AIC and BIC are entirely due to differences in model likelihood. Corresponding versions of the UT and DT models differ in the number of parameters, so these comparisons were based on the AIC and BIC values.

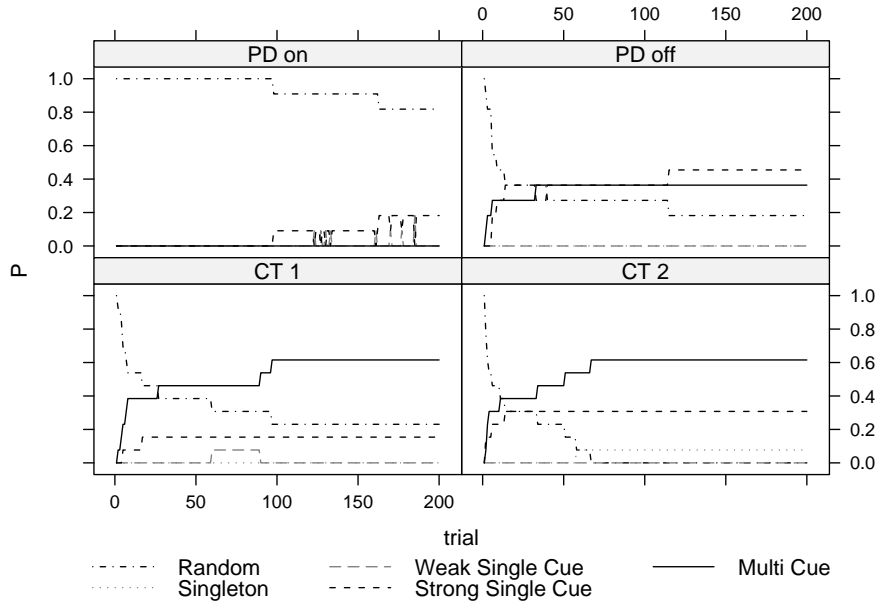


Figure 2. Maximum A Posteriori (MAP) strategy sequences. For the PD on medication condition, these were estimated from the GM/UT model with \mathbf{A}_o and $\lambda_{i,o}$, and for the PD off medication and control conditions (CT1 and CT2, denoting first and second testing occasion respectively), from the CE/DT model with \mathbf{A}_o and $\lambda_{i,o}$.

Cue 1 and Single Cue 4), and “Weak Single Cue” strategies (Single Cue 2 and Single Cue 3). As can be seen there, PD patients on medication appear to have almost exclusively responded according to the Random strategy. PD patients off medication, on the other hand, learned to use both the Strong Single Cue and Multi-Cue strategies. Out of the two PD patients on medication who did switch from Random responding, one switched a total of 3, and one a total of 19 times (these latter switches were mainly from the Single Cue 1 to the Single Cue 3 strategy, and back). In contrast, 9 of the PD patients switched strategies when off medication; the average number of switches (for those that did switch) in this condition was 1.33. The difference in switching behaviour between the two conditions can also be seen in the estimated transition matrices. For the PD patients off medication, the diagonal values a_{ii} , reflecting the probability of strategy continuation, were relatively high ($M = 0.98$, $SD = 0.03$). For the PD patients on medication, the diagonal values were on average much smaller ($M = 0.43$, $SD = 0.32$), indicating more rapid switching behaviour (for those that did switch strategies). In addition, strategy switches of the PD patients off medication were more rational, as the DT model prohibits switches to strategies that decrease expected performance.

The MAP strategy sequences are the most probable strategy sequences given a participant’s responses and the model parameters. This “best guess” of participants’ strategies does not reflect the level of uncertainty associated with these assignments, and as such may miss subtle changes in the evidence for the use of particular strategies. Such subtle changes can be inferred from the strategies’ posterior probabilities ($P(S_t|R_{1:T})$, see

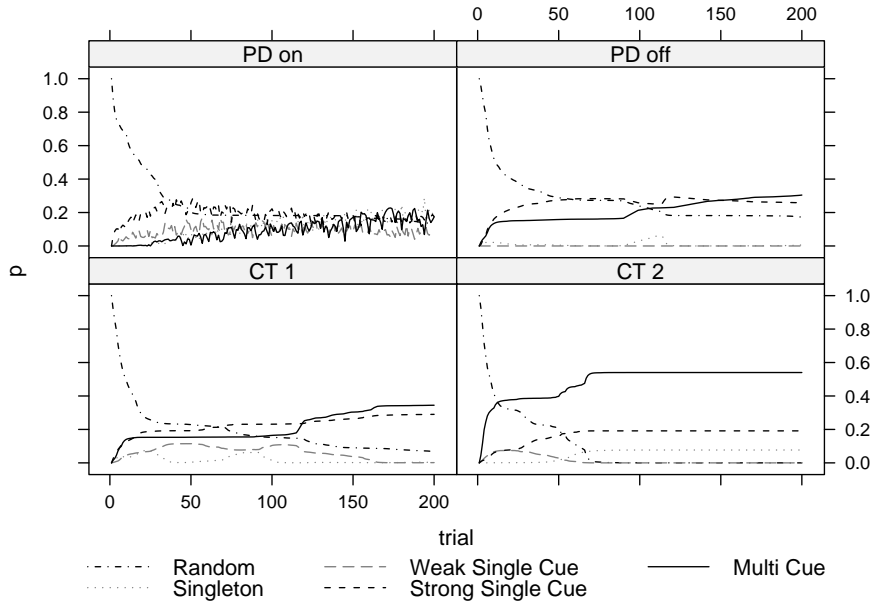


Figure 3. Posterior state probabilities. For the PD on medication condition, these were estimated from the GM/UT model with \mathbf{A}_o and $\lambda_{i,o}$, and for the PD off medication and control conditions (CT1 and CT2, denoting first and second testing occasion respectively), from the CE/DT model with \mathbf{A}_o and $\lambda_{i,o}$.

Appendix) depicted in Figure 3. As can be seen there, the posterior probabilities mainly reproduce the pattern of the MAP strategy sequences, but for the PD on medication condition, they do indicate a slow increase for the Multi-Cue strategy, which was not evident in the MAP sequences. As such, there is some indication that, given more training, PD patients on medication might be able to learn the Multi-Cue strategy.

5.2. Control participants

For the control group, the AIC indicated that the overall best fitting model was the CE/DT model with identical transition probabilities (\mathbf{A}), but differing response consistency between individuals and test occasions ($\lambda_{i,o}$). The BIC selected as overall best fitting the same CE/DT model (with \mathbf{A}), but with identical response consistency between individuals and occasions (λ). Pairwise comparisons confirmed that the CE models were preferred to the GM models, $Z = 3.41$, $p < .001$. In addition, both the AIC and BIC values indicate that the DT models fitted better than the UT models, $Z = 2.59$, $p < .01$ and $Z = 3.52$, $p < .001$, for the AIC and BIC values respectively.

For the CE/DT model, likelihood ratio tests showed that the transition probabilities differed between testing occasions, $\chi^2(12) = 21.63$, $p = .04$. Furthermore, there was clear evidence of individual variability in response consistency, $\chi^2(24) = 82.00$, $p < .001$, and that individual response consistency varied with testing occasion, $\chi^2(13) = 67.30$, $p < .001$. Note that the evidence for different transition probabilities is relatively weak compared to that for varying response consistency. Inspection of the estimated

transition matrices showed little difference, but participants may have been slightly less variable in their strategy use (i.e., fewer strategy switches) on the second testing occasion. Inspection of the response consistency estimates showed that responses were on average more consistent on the second (median $p = 0.88$) than on the first (median $p = 0.86$) test occasion, but this difference was not significant (Wilcoxon signed rank test, $V = 28$, $p = 0.24$).

Comparing the MAP strategy sequences between the two testing occasions (Figure 2), we see that participants on both occasions learned to use the Multi-Cue strategy. However, on the first testing occasion, the Random strategy was also quite prevalent. On the second testing occasion, this latter strategy was overtaken by the Strong Single Cue strategies. The mean number of strategy switches (for those participants who showed at least one switch), was 4.56 and 1.15 on the first and second testing occasion respectively. The relatively high average on the first occasion was mainly due to two participants, one of whom showed a large number of switches between the two Weak Single Cue strategies before settling on the Multi-Cue strategy, and the other switching between the two Strong Single Cue strategies.

5.3. Comparing PD patients and controls

The results above indicate that the PD patients off medication followed a similar strategy switch process as the controls (for both groups, the CE/DT model fitted best). When on medication, the PD patients appeared to use the GM strategies, and switches between these strategies were not always in the direction of increased performance. To further investigate the similarity between the PD patients off medication and controls, and the difference between these conditions and the PD patients on medication, we fitted a number of additional versions of the models. In particular, we were interested in whether PD patients off medication followed a similar strategy switch process – as reflected by the transition matrix \mathbf{A} – in addition to relying on the same CE strategies. Based on the previous results, we took response consistency to vary over individuals and occasions (i.e., $\lambda_{i,o}$). We then investigated different restrictions on the transition matrix. In particular, we fitted a version in which \mathbf{A} was identical for all conditions, and a version in which we took \mathbf{A} to be different in the PD on medication condition, but identical in all other conditions. For comparison, we also fitted a version in which \mathbf{A} was taken to be different for the PD off medication condition, but identical for all other conditions.

The fit measures of these additional models, together with aggregated measures of the models fitted previously, are given in Table 4. In this table, we can see that the AIC selected as the best fitting model the CE/UT model with the transition matrix differing between the PD on medication condition and the other conditions. As earlier, the BIC selected the simplest GM/DT model, with an identical transition matrix for all conditions. For the CE/UT model, the likelihood ratio tests indicate that the transition probabilities were identical between the PD off medication and the two control groups, $\chi^2(84) = 97.98$, $p = .14$. In contrast, the transition probabilities were not identical between the PD on medication condition and the two control conditions, $\chi^2(84) = 144.25$, $p < .001$. In addition, the version which assumes identical transition probabilities for the PD conditions, different to those of the control conditions, was also rejected, $\chi^2(84) = 122.76$, $p < .001$. These analyses thus show that, when off medication, PD patients follow

Table 4

Fit measures of the strategy switch models fitted to the overall data. Best fitting models in bold.

A (id)					Constant Error (CE)			Generalized Matching (GM)		
PD on	PD off	CT 1	CT 2	# par	log L	AIC	BIC	log L	AIC	BIC
Undirected Transition (UT)										
1	2	3	4	216	-4822.39	10076.78	11625.40	-4842.63	10117.26	11665.88
1	2	3	3	174	-4830.48	10008.95	11256.45	-4853.18	10054.37	11301.86
1	1	2	2	132	-4858.49	9980.98	10927.36	-4882.67	10029.35	10975.72
1	2	2	2	132	-4846.10	9956.20	10902.57	-4859.28	9982.55	10928.93
1	2	1	1	132	-4869.24	10002.47	10948.85	-4896.93	10057.86	11004.23
1	1	1	1	90	-4901.31	9982.62	10627.88	-4910.51	10001.03	10646.29
Directed Transition (DT)										
1	2	3	4	96	-4922.60	10037.20	10725.47	-4911.26	10014.51	10702.79
1	2	3	3	84	-4933.41	10034.83	10637.07	-4923.44	10014.89	10617.13
1	1	2	2	72	-4930.47	10004.93	10521.14	-4930.36	10004.71	10520.92
1	2	2	2	72	-4937.19	10018.38	10534.58	-4937.62	10019.24	10535.44
1	2	1	1	72	-4946.01	10036.02	10552.22	-4943.51	10031.02	10547.23
1	1	1	1	59	-4948.02	10016.03	10446.20	-4947.32	10014.64	10444.81

Note: values under **A** (id) indicate equality constraints. E.g., for the models on the second row, there are three unique transition matrices: one for the PD on medication (id = 1), one for the PD off medication (id = 2), and one for the two control conditions (id=3).

a similar strategy switch process to matched controls. This process is different when the PD patients are on medication. The estimated transition probabilities are depicted in Figure 4. As can be seen there, the different strategies are more “connected” in the control and PD off medication conditions. The transition probabilities for the PD on medication condition show two clusters, and starting with the Random strategy, the probability of switching to the Multi-Cue strategy is very small. In addition to showing a similar strategy switch process as controls, PD patients off medication also did not differ in terms of response consistency: median $p = 0.91$ (PD off) vs median $p = 0.94$ (controls), Wilcoxon $W = 122$, $p = 0.49$. The median consistency for the PD patients on medication was 0.88, which was significantly different from the consistency in the control conditions, $W = 82$, $p < .05$, and marginally different from the consistency in the control and PD off medication conditions combined, $W = 125$, $p = .054$.

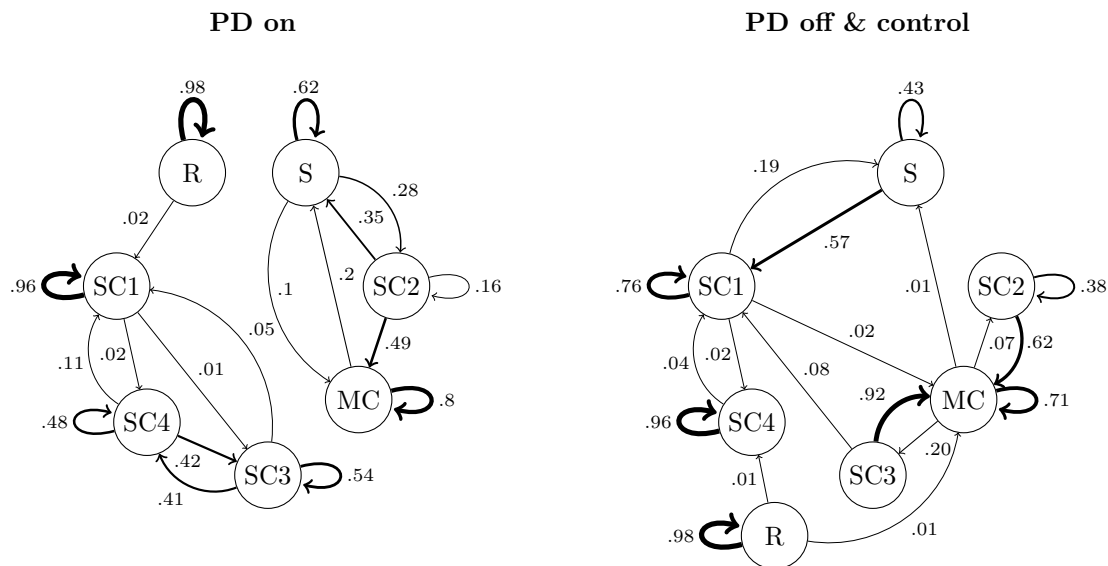


Figure 4. Transition probabilities for the CE/UT model. The PD off medication and control conditions were assumed to have identical transition probabilities. Transition probabilities for the PD on medication group are depicted on the left hand side; transition probabilities for the other conditions on the right hand side. Only transition probabilities $\geq .01$ are depicted. Thicker arrows indicate higher probabilities. R = Random, S = Singleton, SC1-SC4 = Single Cue 1-4, MC = Multi-Cue.

6. Discussion

We presented a general hidden Markov model for rule-based response behaviour in probabilistic category learning tasks. Two types of strategy switch models were described. The Constant Error model, following Gluck et al. (2002) and Meeter et al. (2006), is a “classical” rule model in which participants responded according to an easily verbalisable, deterministic rule (although the Multi-Cue strategy requires a rather lengthy description). The Generalized Matching model takes the same strategies, but responses are assumed to be made by a form of matching to the conditional category probabilities. In the Directed Transition version of these models, we restricted strategy switches to be rational, in the sense that a newly adopted strategy should not decrease the probability of a correct response. In the Undirected Transition version, no such restriction was placed.

We applied the models to the data of Jahanshahi et al. (2009). Their data showed that PD patients on medication were impaired on the WPT, while the same PD patients off medication were not. Our models offered useful results with regards to the role of L-dopa in PCL impairment in Parkinson’s disease. In particular, we showed that PD patients on medication responded more in line with the Generalized Matching model, while PD patients off medication and controls were better modelled by the Constant Error model. In addition, we showed that strategy switches of PD patients off medication and controls were rational, in the sense just described, while those of PD patients on medication were not directly rational. Finally, we showed that PD patients followed a different pattern of strategy switches than controls when on medication, but not when off medication. Both the estimated transition probabilities and MAP strategy sequences indicated that

PD patients on medication mostly relied on a guessing strategy, with some use of simple Single Cue strategies. In contrast, the Multi-Cue strategy was prevalent in PD patients off medication and controls.

6.1. *Learning deficits in PD and L-dopa*

PD patients on medication responded more in accordance with the Generalized Matching than Constant Error version of the strategy switch model. PD patients off medication, like matched controls, responded more in accordance with the Constant Error version. We have argued that the Generalized Matching strategies are more typical of an incremental habit learning system than the Constant Error strategies, which are more indicative of the rule-following behaviour expected from hypothesis testing. The present results thus contradict the hypothesis (e.g., Knowlton et al., 1996) that PCL normally relies on implicit habit learning and that PD patients, due to a dysfunctional habit learning system, must solve the task by declarative strategies. In accordance with later findings (Lagnado et al., 2006; Newell et al., 2007; Price, 2009), our results indicate that PCL typically relies on explicit learning strategies. Moreover, the difference between PD patients on and off medication indicates that these explicit strategies may be impaired by L-dopa, consistent with the L-dopa overdose hypothesis. Indeed, both the PD patients off medication and controls displayed good knowledge of the cue-outcome contingencies, while PD patients on medication were impaired in this respect (see Jahanshahi et al., 2009, for these results). According to the overdose hypothesis, L-dopa can raise dopamine levels beyond optimal in less affected areas such as the ventral striatum and prefrontal cortex. While the dorsal striatum seems to support habitual stimulus-response learning, the ventral striatum and prefrontal cortex appear to support more reflective, goal-directed forms of learning involving cue-outcome representations (Daw et al., 2005; Hampton et al., 2006; O’Doherty et al., 2004). While habitual learning is relatively slow and inflexible, the latter form of learning allows more swift adaptation of behaviour in response to environmental changes. Corresponding to this distinction, evidence suggests that L-dopa impairs performance on tasks which require cognitive flexibility, while enhancing performance on tasks which require cognitive stability (see Cools, 2006). As indicated by the posterior strategy probabilities, the impaired performance of PD patients on medication appears to be due to a general slowness in learning. This could be due to a reliance on slow, and likely impaired, habit learning, although other explanations are also possible.

The overdose hypothesis relates the effects of L-dopa to baseline dopamine levels. As the PD patients had a relatively wide range of disease duration and severity, the effect of L-dopa may have varied between patients. Due to the progressive nature of PD, those with longer duration PD are likely to be affected in more brain areas than those with shorter duration PD. In addition, individuals with long duration PD may respond differently to L-dopa (cycle on and off rapidly). As such, the overdose effect may be particularly pronounced for patients in the early stages of PD. Because the data contained a relatively small sample of PD patients, we did not investigate the effect of disease duration. Future research with larger samples could contrast PD patients in early and late stages and also measure baseline dopamine levels in the relevant areas (i.e., the dorsal and ventral striatum and prefrontal cortex) to investigate whether and how these interact with the effect of L-dopa.

According to the present analysis, a main problem for PD patients on medication was that many did not switch from a random guessing strategy to more advantageous ones. Previous research (e.g., Shohamy et al., 2006) attributed impaired performance of PD patients (on medication) to a failure to switch from the Singleton strategy to more advantageous ones. We found that the Singleton strategy was hardly ever used, which may be due to the inclusion of the Random strategy in our model – as the Singleton strategy is closest to random responding, exclusion of the Random strategy will result in more assignments to the Singleton strategy. As such, assignments of the Singleton strategy in previous studies should be treated with some caution.

While the Random strategy fitted their responses better than the other strategies, PD patients on medication did show an increase in performance during the task (the average performance increased from .60 in block 1 to .65 in block 4). Thus, there seems to be some level of learning which is not captured by the MAP strategy sequences. Indeed, the posterior state probability of the Multi-Cue strategy for PD patients on medication did increase during the task. As such, even PD patients on medication may learn to use the Multi-Cue strategy after sufficient training. In addition, the high prevalence of the Random strategy in the PD on medication group may be due to the lower bound set on response consistency. To increase the discrimination between the strategies, we set a lower bound of $p = .8$ for the Constant Error, and a lower bound of $\lambda = 1$ (pure probability matching) for the Generalized Matching versions. As a result, we may have missed the strategies of participants who “under-shoot” ($\lambda < 1$), or who applied a strategy with $p < .8$.

Undershooting can be expected in the early stages of incremental learning if responses are made by probability matching, and the length of this early stage will depend on the rate of learning. The present analysis is not very sensitive to slow and gradually changing response processes. In other work (Speekenbrink and Shanks, 2008), we have used a dynamic generalized linear model to analyse learning in the WPT. This model has similarities to the hidden Markov model used here, but assumes a (multidimensional) continuous state space, rather than the discrete state space of the strategy switch models, making it more appropriate for a slow, incremental learning mechanism. While certainly interesting, a comparison of this model to the present strategy switch models is beyond the scope of this article. One complication is that maximum likelihood estimation of dynamic generalized linear models is rather difficult, while the MCMC technique we used previously (Speekenbrink and Shanks, 2008) is inadequate for present purposes. We are currently working on a sampling-based method for maximum likelihood estimation of dynamic generalized linear models, which would allow direct comparison with the strategy switch models. We can then more directly assess whether PD patients on medication are indeed best characterized by an incremental “habit” learning process. While we have argued that this form of learning is more consistent with the Generalized Matching strategies, the better fit of the GM model is not sufficient to conclude that PD patients on medication relied on habit learning. An alternative view is that the GM and CE strategies reflect a difference in decision rather than response strategy. Both the GM and CE strategies could derive from the same underlying (incrementally) formed cue-outcome representations (cf. Shohamy et al., 2008). The difference between the strategies then reflects how this information is used. In the Constant Error model, the process is essentially one of maximizing, predicting the most probable outcome, but this maximizing strategy is corrupted by a fixed level of noise. In the Variable Matching

model, the probability of a maximizing response depends directly on the conditional outcome probabilities, indicating a possibly less reflective use of learned information. Future research could investigate this possibility by separating learning and decision processes more clearly. This requires both careful modelling and informative data and we plan to address both in future work.

6.2. *Mathematical models in clinical assessment*

A main contribution of the present study was to develop a statistically well-founded model of strategy switches in PCL. The model was based on the work of Gluck et al. (2002), later extended by Meeter et al. (2006). In comparison to these other models, our strategy switch model has a number of advantages. Firstly, we can obtain maximum likelihood estimates of the strategy transition probabilities. We illustrated that these can show interesting differences between groups. Secondly, group differences can be tested with likelihood ratio tests, comparing a model in which parameters are constrained to be identical between groups to a model in which parameters are allowed to vary between groups. This is a more sensitive method than testing for differences in parameter estimates using post-hoc t-tests or nonparametric equivalents. Thirdly, by fixing transition probabilities, we can assess the rationality of the strategy switch process. Finally, we can estimate maximum a posteriori (MAP) strategy sequences for individual participants. These take into account both the likelihood of a strategy given observations and the strategy transition probabilities, so that they are consistent with fundamental aspects of the model.

The original strategy model has already proven a popular tool to investigate differences between patient groups in PCL (Fera et al., 2005; Shohamy et al., 2004b,a; Wilkinson et al., 2008). This illustrates the large potential of mathematical models in studies with clinical populations. We hope that those interested in strategy use in probabilistic category learning will consider the models described here for future work². The current approach offers a powerful framework to investigate clinically significant differences in PCL. Probabilistic category learning is not just an experimental paradigm – in daily life, people will encounter many situations in which they need to learn to make decisions based on probabilistic feedback. Understanding the precise nature of the cognitive impairments associated with PD, and the effects of medication, is not only interesting for theoretical neuroscience. The finding that one of the current treatments of choice (L-dopa) has adverse effects on learning has important clinical implications because cognitive deficits in PD patients are predictors of poor quality of life in PD patients (Schrag et al., 2000) and their carers (Aarsland et al., 1999), and admission to nursing homes (Aarsland et al., 2000). Hopefully, the present findings will stimulate research into new medication that does not impair cognitive functioning.

² The software used for the current analyses is available from the first author's website via http://www.ucl.ac.uk/psychlangsci/staff/cpb-staff/m_speekenbrink.

7. Acknowledgements

The support of the Economic and Social Research Council (ESRC) is gratefully acknowledged. This research was part of the programme of the ESRC Research Centre for Economic Learning and Social Evolution, University College London (M.S., D.A.L., D.R.S). L.W. was supported by a postdoctoral fellowship from the Parkinson's Disease Society, UK.

The authors would like to thank Mark Gluck and an anonymous reviewer for helpful comments on a previous version of the paper.

Appendix A. Estimation of the strategy switch model

In this appendix, we provide details of the parameter estimation of the strategy switch model, and the determination of the maximum a posteriori state sequences.

A.1. Parameter estimation by EM

To apply the EM algorithm, we treat the latent states as missing observations. Parameters are estimated by iteratively maximising the expected joint likelihood of the parameters given the observations and states. Let $R_{1:T}^{1:N}$ and $S_{1:T}^{1:N}$ denote the set of N repeated series of observations (states) of length T (i.e., N is the number of participants, and T the number of learning trials in the WPT for each participant). From (1) and (2), and noting that we fixed $P(S_1^{1:N})$, the joint log likelihood can be written as

$$\ln P(R_{1:T}^{1:N} = r_{1:T}^{1:N}, S_{1:T}^{1:N} = s_{1:T}^{1:N} | \mathbf{A}, \boldsymbol{\lambda}) = \sum_{i=1}^N \sum_{t=2}^T \ln P(s_t^i | s_{t-1}^i) + \sum_{i=1}^N \sum_{t=1}^T \ln P(r_t | s_t, \lambda_i)$$

This likelihood depends on the unobserved states s_t^i . In the Expectation step, we replace these with their expected values given a set of (initial) parameters $\boldsymbol{\theta}' = (\mathbf{A}, \boldsymbol{\lambda})$ and observations $R_{1:T}^{1:N}$. The expected log likelihood

$$Q(\boldsymbol{\theta}, \boldsymbol{\theta}') = E_{\boldsymbol{\theta}'}(\ln P(R_{1:T}^{1:N}, S_{1:T}^{1:N} | R_{1:T}^{1:N}, \mathbf{A}, \boldsymbol{\lambda}))$$

can be written as

$$Q(\boldsymbol{\theta}, \boldsymbol{\theta}') = \sum_{i=1}^N \sum_{t=2}^T \sum_{j=1}^M \sum_{k=1}^M \xi_t^i(j, k) \ln a_{jk} + \sum_{i=1}^N \sum_{t=1}^T \sum_{j=1}^M \gamma_t^i(j) \ln P(r_t^i | \lambda_i), \quad (\text{A.1})$$

where the expected values $\xi_t^i(j, k) = P(S_t^i = k, S_{t-1}^i = j | R_{1:T}^i, \boldsymbol{\theta}')$ and $\gamma_t^i(j) = P(S_t^i = j | R_{1:T}^i, \boldsymbol{\theta}')$ can be computed effectively by the Forward-Backward algorithm (see e.g., Rabiner, 1989). The Maximisation step consists of the maximisation of (A.1) for $\boldsymbol{\theta}$. As the r.h.s. of (A.1) consists of two separate parts, we can maximise separately for \mathbf{A} and $\boldsymbol{\lambda}$. The maximising values of a_{jk} are (e.g., Rabiner, 1989)

$$\hat{a}_{jk} = \frac{1}{N(T-1)} \sum_{i=1}^N \sum_{t=2}^T \frac{\xi_t^i(j, k)}{\gamma_{t-1}^i(j)}. \quad (\text{A.2})$$

Fixing certain elements of \mathbf{A} to 0, as in the DT version, does not affect the estimation of the other elements. When elements a_{jk} are assumed identical, we simply extend the summation in (A.2) to include all those elements, changing the denominator $N(T - 1)$ accordingly.

To estimate $\boldsymbol{\lambda}$, we note that the term containing this parameter has the form of a weighted likelihood, where $\gamma_t^i(j)$ can be interpreted as the number of replications of r_t^i . Hence, we can rely on standard ML estimates of the logistic regression coefficients λ_i , using the values $\gamma_t^i(j)$ as “case weights” (e.g., Agresti, 2002; McCullagh and Nelder, 1983)³. The resulting parameter estimates are then used in a new Expectation step, followed by another Maximisation step, until convergence. Under mild regularity conditions, it can be shown that the EM algorithm converges to a (local) maximum of the likelihood (Dempster et al., 1977).

A.2. Maximum A Posteriori state sequences

Note that, at the ML estimate of the parameters, the expectations $\gamma_t^i(j)$ are the posterior state probabilities $P(S_t^i = j | R_{1:T}^i)$, and can be used to determine point-wise Maximum A Posteriori (MAP) states. However, the resulting point-wise MAP state *sequence* might be quite different from the overall MAP state sequence (for instance, it can contain state transitions with probability 0). The latter can be computed by the Viterbi algorithm (e.g., Rabiner, 1989). For completeness, we give the relevant details below.

The Viterbi algorithm computes, for each individual i , the maximum a posteriori state sequence

$$s^{i*} \equiv \arg \max_{s_{1:T}} P(S_{1:T}^i = s_{1:T} | R_{1:T}^i).$$

For readability, we will omit the superscript i in the remainder. The Viterbi algorithm works with two variables. The first, $\delta_t(j)$, keeps track of the maximum probability of reaching state j at trial t for a single state sequence, i.e.

$$\delta_t(j) \equiv \max_{s_{1:(t-1)}} P(S_{1:(t-1)} = s_{1:(t-1)}, S_t = j | R_{1:t}).$$

The second variable keeps track of which preceding state sequence gives state j at t the highest probability, i.e.

$$\psi_t(j) \equiv \arg \max_k \delta_{t-1}(k) a_{kj}.$$

We initialize the variables at

$$\begin{aligned} \delta_1(j) &= P(S_1 = j)P(R_1 | S_1 = j) \\ \psi_1^i(j) &= 0. \end{aligned}$$

and for $t = 2, \dots, T$, we have

³ To be more specific, we replicate each observation r_t^i a total of 6 times, once for each of the states besides the random state (which offers no information regarding λ_i). For the j -th replication (corresponding to the j -th state), we used $v_j(\mathbf{x}_t)$ as a predictor variable and $\gamma_t^i(j)$ as a case weight. All these replications were used to obtain the maximum likelihood estimate of λ from a single GLM, using the “glm.fit” function in R (R Development Core Team, 2006).

$$\delta_t(j) = \max_k \delta_{t-1}(k) a_{kj} P(R_t | S_t = j)$$

$$\psi_t(j) = \arg \max_k \delta_{t-1}(k) a_{kj}.$$

Once the values $\delta_t(j)$ and $\psi_t(j)$ have been computed, determination of the MAP state sequence is straightforward. Clearly,

$$s_T^* = \arg \max_j \delta_T(j).$$

For $t = T - 1, \dots, 1$, we can then proceed as

$$s_t^* = \psi_{t+1}(s_{t+1}^*).$$

Note that, especially for long sequences, it will be numerically better to work with $\delta'_t(j) = \log \delta_t(j)$. This does not affect the determination of ψ_t .

References

- Aarsland, D., Larsen, J. P., Karlsen, K., Lim, N. G., Tandberg, E., 1999. Mental symptoms in parkinson's disease are important contributors to caregiver distress. *International Journal of Geriatric Psychiatry* 14, 866–874.
- Aarsland, D., Larsen, J. P., Tandberg, E., Laake, K., 2000. Predictors of nursing home placement in parkinson's disease: a population-based, prospective study. *Journal of the American Geriatrics Society* 48, 938–942.
- Aarsland, D., Zaccai, J., Brayne, C., 2005. A systematic review of prevalence studies of dementia in Parkinson's disease. *Movement Disorders* 20, 1255–1263.
- Agresti, A., 2002. *Categorical data analysis*, 2nd Edition. Wiley, Hoboken, NJ.
- [Akaike73 Akaike, H., 1992[1973]. Information theory and an extension of the maximum likelihood principle. In: Kotz, S., Johnson, K. L. (Eds.), *Breakthroughs in statistics*. Vol 1. Springer-Verlag, London, pp. 610–624.
- 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., Maddox, W. T., 1993. Relations between prototype, exemplar, and decision bound models of categorization. *Journal of Mathematical Psychology* 37, 372–400.
- Ashby, F. G., Maddox, W. T., 2005. Human category learning. *Annual Review of Psychology* 56, 149–178.
- Ashby, F. G., Noble, S., Filoteo, J. V., Waldron, E. M., Ell, S. W., 2003. Category learning deficits in parkinson's disease. *Neuropsychology* 17, 115–124.
- Baum, W. M., 1974. On two types of deviation from the matching law: Bias and undermatching. *Journal of the Experimental Analysis of Behavior* 22, 231–242.
- Baum, W. M., 1979. Matching, undermatching, and overmatching in studies of choice. *Journal of the Experimental Analysis of Behavior* 32, 269–281.
- Cools, R., 2006. Dopaminergic modulation of cognitive function-implications for l-dopa treatment in parkinson's disease. *Neuroscience and Biobehavioral Reviews* 30, 1–23. URL <http://dx.doi.org/10.1016/j.neubiorev.2005.03.024>
- Cools, R., Barker, R. A., Sahakian, B. J., Robbins, T. W., 2001a. Enhanced or impaired cognitive function in parkinson's disease as a function of dopaminergic medication and task demands. *Cerebral Cortex* 11, 1136–1143.

- Cools, R., Barker, R. A., Sahakian, B. J., Robbins, T. W., 2001b. Mechanisms of cognitive set flexibility in parkinsons disease. *Brain* 124, 2503–2512.
- 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.
- Dempster, A. P., Laird, N. M., Rubin, D. B., 1977. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society, Series B* 39, 1–38.
- Dobson, A. J., 2002. An introduction to generalized linear models, 2nd Edition. Chapman & Hall, Boca Raton.
- Eldridge, L. L., Masterman, D., Knowlton, B. J., 2002. Intact implicit habit learning in Alzheimer’s disease. *Behavioral Neuroscience* 116, 722–726.
- Fera, F., Weickert, T. W., Goldberg, T. E., Tessitore, A., Hariri, A., Das, S., Lee, S., Zolnick, B., Meeter, M., Myers, C. E., Gluck, M. A., Weinberger, D. R., Mattay, V. S., 2005. Neural mechanisms underlying probabilistic category learning in normal aging. *Journal of Neuroscience* 25, 11340–11348.
- Filoteo, J. V., Maddox, W. T., 1999. Quantitative modeling of visual attention processes in patients with Parkinson’s disease: Effects of stimulus integrality on selective attention and dimensional integration. *Neuropsychology* 13, 206–222.
- 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.
- Filoteo, J. V., Maddox, W. T., Salmon, D. P., Song, D. D., Mar 2005. Information-integration category learning in patients with striatal dysfunction. *Neuropsychology* 19 (2), 212–222.
- 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.
- Gabrieli, J. D., 1998. Cognitive neuroscience of human memory. *Annual Review of Psychology* 49, 87–115.
- 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 & Memory* 9, 408–418.
- Gotham, A. M., Brown, R. G., Marsden, C. D., 1988. ‘Frontal’ cognitive function in patients with Parkinson’s disease ‘on’ and ‘off’ levodopa. *Brain* 111, 299–321.
- Hampton, A. N., Bossaerts, P., O’Doherty, J. P., 2006. The role of the ventromedial prefrontal cortex in abstract state-based inference during decision making in humans. *Journal of Neuroscience* 26, 8360–8367.
- Hoehn, M., Yahr, M., 1967. Parkinsonism: Onset, progression, and mortality. *Neurology* 17, 427–442.
- Jahanshahi, M., Wilkinson, L., Gahir, H., Dharminda, A., Lagnado, D. A., 2009. Medication impairs probabilistic classification learning in Parkinson’s disease. Manuscript submitted for publication.
- Kelley, H., Friedman, D., 2002. Learning to forecast price. *Economic Inquiry* 40, 556–573.
- Knowlton, B. J., Mangels, J. A., Squire, L. R., 1996. A neostriatal habit learning system in humans. *Science* 273, 1399–1402.
- Knowlton, B. J., Squire, L. R., Gluck, M. A., 1994. Probabilistic classification learning in amnesia. *Learning & Memory* 1, 106–120.

- Kollins, S., Newland, M., Critchfield, T., 1997. Human sensitivity to reinforcement in operant choice: How much do consequences matter? *Psychonomic Bulletin & Review* 4, 208–220.
- Kruschke, J. K., Johansen, M. K., 1999. A model of probabilistic category learning. *Journal of Experimental Psychology: Learning, Memory and Cognition* 25, 1083–1119.
- Lagnado, D. A., Newell, B. R., Kahan, S., Shanks, D. R., 2006. Insight and strategy in multiple cue learning. *Journal of Experimental Psychology: General* 135, 162–183.
- Luce, D. R., 1959. *Individual choice behavior*. Wiley, New York.
- Maddox, W. T., Aparicio, P., Marchant, N. L., Ivry, R. B., May 2005. Rule-based category learning is impaired in patients with parkinson’s disease but not in patients with cerebellar disorders. *Journal of Cognitive Neuroscience* 17 (5), 707–723.
- McCullagh, P., Nelder, J. A., 1983. *Generalized linear models*. Chapman & Hall, London.
- Meeter, M., Myers, C. E., Shohamy, D., Hopkins, R. O., Gluck, M. A., 2006. Strategies in probabilistic categorization: results from a new way of analyzing performance. *Learning & Memory* 13 (2), 230–239.
- Meeter, M., Radics, G., Myers, C., Gluck, M., Hopkins, R., 2008. Probabilistic categorization: how do normal participants and amnesic patients do it? *Neuroscience and Biobehavioral Reviews* 32, 237–48.
- Moody, T. D., Bookheimer, S. Y., Vanek, Z., Knowlton, B. J., 2004. An implicit learning task activates medial temporal lobe in patients with Parkinson’s disease. *Behavioral Neuroscience* 118, 438–442.
- Newell, B. R., Lagnado, D. A., Shanks, D. R., 2007. Challenging the role of implicit processes in probabilistic category learning. *Psychonomic Bulletin & Review* 14, 505–511.
- Nieoullon, A., 2002. Dopamine and the regulation of cognition and attention. *Progress in Neurobiology* 67, 53–83.
- O’Doherty, J., Dayan, P., Schultz, J., Deichmann, R., Friston, K., Dolan, R. J., 2004. Dissociable roles of ventral and dorsal striatum in instrumental conditioning. *Science* 304, 452–454.
- Owen, A. M., Iddon, J. L., Hodges, J. R., Summers, B. A., Robbins, T. W., 1997. Spatial and non-spatial working memory at different stages of Parkinson’s disease. *Neuropsychologia* 35, 519–532.
- Poldrack, R. A., Clark, J., Par-Blagoev, E. J., Shohamy, D., Moyano, J. C., Myers, C., Gluck, M. A., 2001. Interactive memory systems in the human brain. *Nature* 414, 546–550.
- Poldrack, R. A., Rodriguez, P., 2004. How do memory systems interact? Evidence from human classification learning. *Neurobiology of Learning and Memory* 82, 324–332.
- Price, A. L., 2005. Cortico-striatal contributions to category learning: Dissociating the verbal and implicit systems. *Behavioral Neuroscience* 119, 1438–1447.
- Price, A. L., 2009. Distinguishing the contributions of implicit and explicit processes to performance of the weather prediction task. *Memory & Cognition* 37, 210–222.
- R Development Core Team, 2006. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, ISBN 3-900051-07-0. URL <http://www.R-project.org>
- Rabiner, L. R., 1989. A tutorial on hidden Markov models and selected applications in speech recognition. *Proceedings of the IEEE* 77, 257–286.
- Rescorla, R. A., Wagner, A. R., 1972. A theory of Pavlovian conditioning: Variations in

- the effectiveness of reinforcement and nonreinforcement. In: Black, A. H., Prokasy, W. F. (Eds.), *Classical conditioning II: Current theory and research*. Appleton-Century-Cross, New York, pp. 64–99.
- Sage, J. R., Anagnostaras, S. G., Mitchell, S., Bronstein, J. M., Salles, A. D., Masterman, D., Knowlton, B. J., 2003. Analysis of probabilistic classification learning in patients with Parkinson’s disease before and after pallidotomy surgery. *Learning & Memory* 10, 226–236.
- Saint-Cyr, J., Taylor, A., Lang, A., 1988. Procedural learning and neostriatal dysfunction in man. *Brain* 111, 941–959.
- Schrag, A., Jahanshahi, M., Quinn, N., 2000. What contributes to quality of life in patients with parkinson’s disease? *Journal of Neurology, Neurosurgery, and Psychiatry* 69, 308–312.
- Schultz, W., 2002. Getting formal with dopamine and reward. *Neuron* 36, 241–263.
- Schultz, W., Dayan, P., Montague, P. R., 1997. A neural substrate of prediction and reward. *Science* 275, 1593–1599.
- Schwarz, G., 1978. Estimating the dimension of a model. *Annals of Statistics* 6, 461–464.
- Shohamy, D., Meyers, 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. L-dopa impairs learning, but spares generalization, in parkinson’s disease. *Neuropsychologia* 44, 774–784.
- Shohamy, D., Myers, C. E., Grossman, S., Sage, J., Gluck, M. A., Poldrack, R. A., 2004a. Cortico-striatal contributions to feedback-based learning: converging data from neuroimaging and neuropsychology. *Brain* 127, 1–9.
- Shohamy, D., Myers, C. E., Onlaor, S., Gluck, M. A., 2004b. Role of the basal ganglia in category learning: How do patients with Parkinson’s disease learn? *Behavioral Neuroscience* 118, 676–686.
- Speekenbrink, M., Channon, S., Shanks, D. R., 2008. Learning strategies in amnesia. *Neuroscience and Biobehavioral Reviews* 32, 292–310.
- Speekenbrink, M., Shanks, D. R., 2008. Through the looking glass: A dynamic lens model approach to learning in MCPL tasks. In: Chater, N., Oaksford, M. (Eds.), *The probabilistic mind: Prospects for rational models of cognition*. Oxford University Press, Oxford, pp. 409–429.
- Swainson, R., Rogers, R. D., Sahakian, B. J., Summers, B. A., Polkey, C. E., Robbins, T. W., 2000. Probabilistic learning and reversal deficits in patients with parkinson’s disease or frontal or temporal lobe lesions: possible adverse effects of dopaminergic medication. *Neuropsychologia* 38, 596–612.
- Wickens, T. D., 1982. *Models for behavior: Stochastic processes in psychology*. W. H. Freeman, San Francisco.
- Wilkinson, L., Lagnado, D. A., Quallo, M., Jahanshahi, M., 2008. The effect of corrective feedback on non-motor probabilistic classification learning in parkinson’s disease. *Neuropsychologia* 46, 2683–2695.
- Witt, K., Nushman, A., Deuschl, G., 2002. Dissociation of habit-learning in Parkinson’s and cerebellar disease. *Journal of Cognitive Neuroscience* 14, 493–499.