

RUNNING HEAD: Probabilistic category learning

Challenging the Role of Implicit Processes in Probabilistic Category Learning

Ben R. Newell

University of New South Wales

David A. Lagnado and David R. Shanks

University College London

Address correspondence to:

Ben R. Newell

School of Psychology

University of New South Wales

Sydney 2052

AUSTRALIA

Tel: +61 2 9385 1606

Fax: +61 2 9385 3641

ben.newell@unsw.edu.au

Abstract

Considerable interest in the hypothesis that different cognitive tasks recruit qualitatively distinct processing systems has led to the proposal of separate explicit (declarative) and implicit (procedural) systems. A popular probabilistic category learning task known as the “Weather Prediction Task” is said to be ideally suited to examine this distinction because its two versions – ‘observation’ and ‘feedback’ – are claimed to recruit the declarative and procedural systems respectively. In two experiments we found results that were inconsistent with this interpretation. In Experiment 1 a concurrent memory task had a greater detrimental effect on the putatively implicitly mediated (feedback) version than on the explicit (observation) version. In Experiment 2 participants displayed comparable and accurate insight into the task and their judgment processes in both the feedback and observation versions. These findings have important implications for the study of probabilistic category learning in both normal and patient populations.

In recent years there has been considerable interest in the hypothesis that different cognitive tasks recruit qualitatively distinct processing systems (e.g. Ashby, Alfonso-Reese, Turken & Waldron, 1998; Gabrieli, 1998; Knowlton, Squire & Gluck, 1994; Reber, Knowlton & Squire, 1996; Squire, 2004). The major distinction drawn by proponents of this multiple-systems view is between an explicit or declarative system that requires awareness and involves analytic processing, and an implicit or procedural system that operates in the absence of conscious awareness and is non-analytic. In addition to the distinction in terms of operational properties, some researchers suggest that specific neuroanatomical regions are differentially involved in mediating the two systems. One popular view is that the basal ganglia are involved in procedural learning, whereas the medial temporal lobes are involved in declarative learning (e.g. Ashby et al., 1998; Poldrack, Clark, Pare-Blagoev, Shohamy, Creso Moyano, & Myers, 2001; Shohamy, Myers, Onlaor & Gluck, 2004; Shohamy, Myers, Grossman, Sage, Gluck & Poldrack, 2004; Squire, 2004).

A key piece of evidence for this proposed dissociation is the difference in performance between tasks that are learned via trial-by-trial feedback and those which are learned in an ‘observational’ manner with no feedback (e.g. Ashby, Maddox & Bohil, 2002; Poldrack et al., 2001; Shohamy et al., 2004a). Several authors have claimed that the procedural system is recruited when tasks are complex and learned incrementally via feedback. The feedback is crucial to engage the operation of a reward-related learning system thought to be mediated via dopamine neurons in the basal ganglia (Ashby et al., 1998; Schultz, Dayan, & Montague, 1997). Observation tasks are those in which, typically, stimuli are shown together with the correct outcome, no behavioral response is required and no feedback is provided. These tasks fail to engage the procedural system because there is no ‘surprising’ reward associated with learning. With such observational learning tasks, it is argued that performance on complex tasks is either impaired (Ashby et al., 2002) or a

qualitatively different and neuroanatomically distinct system – (i.e. the declarative system) – takes over learning (Poldrack et al., 2001; Shohamy et al., 2004a) as the task requires a more pure form of memorization.

Some of the evidence for these different patterns of performance comes from neuropsychological studies. Studies with Parkinson's patients are of particular relevance because these patients suffer from a profound loss of dopamine-containing neurons in the substantia nigra. The loss of these neurons causes a decrease in striatal dopamine and disruption of basal ganglia function (Shohamy, Myers, Onalor, & Gluck, 2004). Filoteo, Maddox, Salmon, and Song (2004) demonstrated that Parkinson's patients were impaired at learning complex 'information integration' tasks; the learning of which is thought to rely on the procedural system (though see Ashby, Ell, & Waldron, 2003). Similarly, Shohamy et al. (2004a) demonstrated that Parkinson's patients were impaired in learning a feedback version of a probabilistic categorization task, but unimpaired, relative to controls, on an observation version of the same task.

There are also several claims for differential involvement of the two systems in non-patient groups. A study by Ashby et al. (2002) comparing performance on observation and feedback versions of information integration tasks showed impairments in the observation version in normal participants. Poldrack et al. (2001) demonstrated equivalent performance in normals on feedback and observation versions of a probabilistic categorization task but presented neuroimaging data suggesting differential engagement of the basal ganglia and the medial temporal lobe in the different versions.

Presenting this dichotomous view of cognition, in which different tasks are claimed to recruit different systems, naturally invites simple experimental manipulations hypothesized to have differential effects on the two putative systems (see Ashby & Maddox, 2005 and Maddox & Ashby, 2004 for reviews of a comprehensive research program with this aim). In

this paper, we draw on this simple strategy and scrutinize the processes underlying performance in a task that has been one of the primary sources of data for the proposed dissociation between the procedural and declarative system. The task is a prototypical probabilistic category learning task known as the weather prediction task (Knowlton et al., 1994). The task has been used in numerous investigations with a variety of populations (e.g. unimpaired individuals: Gluck, Shohamy & Myers, 2002; Lagnado, Newell, Kahan, & Shanks, 2005; Alzheimer's patients: Eldridge, Masterman & Knowlton, 2002; Parkinson's patients: Moody, Bookheimer, Vanek & Knowlton, 2004; Shohamy et al. 2004a; Shohamy et al. 2004b; amnesic individuals: Hopkins, Myers, Shohamy, Grossman & Gluck, 2004; Reber et al., 1996; and schizophrenics: Keri, Kelemen, Szekeres, Bagoczky, Erdelyi, Antal, Benedek & Janka, 2000) and has therefore been highly influential in the development of theoretical models of dissociable learning and memory systems (Squire, 2004).

In the weather prediction task people learn to predict a binary outcome (rainy or fine weather) on the basis of four binary cues (four distinct tarot cards). Each card is associated with the outcome with a different probability, and these combine to determine the probability of the outcome for any particular pattern of cards. The trials presented in an experiment are made up of a representative sampling of possible card combinations. In the feedback version of the task, on each trial participants see a specific pattern of cards, predict the weather, and receive corrective feedback as to the correct outcome. This enables them to gradually learn the cue-outcome relations and thus improve the accuracy of their predictions. This feedback version is claimed to be mediated via the procedural system (Gluck et al., 2002; Knowlton et al., 1994; Poldrack et al., 2001; Reber et al., 1996; Shohamy et al., 2004a; Shohamy et al., 2004b).

In the observation version of the task, on each trial participants are presented with both the cards and the outcome simultaneously. Crucially, this paired-associate arrangement

does not rely on trial-by-trial feedback and so learning the task does not (perhaps cannot, cf. Ashby et al., 1998) recruit the procedural system. Rather, the declarative system is thought to be primarily responsible for learning (Poldrack et al., 2001; Shohamy et al., 2004a).

Given this purported distinction between the two versions of the task, clear predictions can be made about the ways in which various experimental manipulations should affect performance in the observation and feedback versions. In Experiment 1 we make the simple prediction that if performance in the feedback version is mediated by an implicit procedural system (that is neuroanatomically distinct from the explicit system) then it should be less affected by placing additional demands on working memory. In contrast, the observation version, thought to be primarily under declarative control, should be detrimentally affected by an additional load on working memory.

The logic of this manipulation is the same as that used by Waldron and Ashby (2001) in their investigation of the effects of concurrent working memory tasks on performance in information-integration and rule-based tasks. Although the feedback/observation distinction does not map directly on to the information-integration/rule-based one (see the General Discussion for elaboration of this point) the logic of the manipulation is the same. In their study Waldron and Ashby found that execution of a concurrent memory task (a numerical Stroop task) led to greater decrements in performance on a rule based task than an information-integration task. This finding lends support to their argument that the tasks are respectively mediated by explicit (working memory dependent) and implicit (working memory independent) systems.

If we were to find such a pattern of results with the two versions of the weather prediction task then it would lend support to the idea that separate systems underlie performance in the feedback and observation versions. If on the other hand we were to find a general decrement across both tasks, or indeed a greater decrement in the feedback version

than the observation version then this would question the current interpretation of performance in the weather prediction task.

We should stress that part of our motivation for conducting the current experiments was recent research that has already questioned the characterization of the feedback and observation versions of the WEATHER PREDICTION task as being *a priori* procedural and declarative (e.g. Ashby & Maddox, 2005; Gluck et al., 2002; Lagnado et al., 2005; Shohamy et al., 2004a). Part of the problem with the task is that although it was designed as one that would be reliant on the procedural learning system, there are a number of ways in which a participant might attempt to ‘solve’ the task.

Gluck et al. (2002) identified three types of strategies that appeared to account for the performance of most participants trying to solve the feedback version of the task. These were: (1) a multi-cue strategy in which participants learn about all four cards, and base their predictions on some integration of this information; (2) a singleton strategy in which they just learn about the cue patterns with a single card (and guess when more than one card is present); (3) a one-cue strategy in which they focus on just one card, and base their predictions on the presence or absence of this card. The structure of the task environment is such that these simpler strategies can be almost as successful as the more complex ones.

Interestingly many participants in Gluck et al.’s experiments appeared to use the simpler, and potentially ‘explicit’, singleton strategy. However, Gluck et al. (2002) maintained that even this simple strategy could have been the product of an implicit system as the strategy identified through participants’ behavioral responses was typically unrelated to their verbal reports of the strategies they had adopted. Strategy analyses of the observation version of the task are as yet inconclusive, but Shohamy et al. (2004a) noted that aspects of their data from the observational version were inconsistent with the use of the explicit analytic strategies that would have been expected if the task was under declarative control.

Perhaps even more importantly, Lagnado et al. (2005) recently provided strong evidence that the feedback version of the task involves explicit rather than implicit processes. Lagnado et al. (2005) elicited trial-by-trial ratings of participants' reliance on each cue and demonstrated a clear divergence in the ratings given to weak and strong predictors of the outcome. This self-insight was coupled with excellent task knowledge demonstrated through an ability to judge accurately the probability of an outcome given the presence of a particular cue. In Experiment 2 we use these more sensitive and dynamic measures of insight to investigate the simple prediction that insight and task knowledge in the 'declarative' observation version should be superior to that in the 'procedural' feedback version. Again, if we find this to be the case it will lend support to the suggestion that the different versions of the task are mediated by different systems. If we find no difference or superior insight in the feedback version then again the current interpretation of the weather prediction task will be questioned.

In the following experiments we capitalize on these recent advances in analyzing performance and insight in the weather prediction task. Our aim is to test the two simple predictions that arise from the multiple-systems view: 1) that placing additional demands on working memory should be more deleterious to performance in the observation version than the feedback version and 2) that insight and task knowledge should be superior following observation training than feedback training. Finding support for these two predictions would help to justify the extensive use of the weather prediction task in cognitive neuroscience as a 'tool' for demonstrating the operation of the hypothesized procedural and declarative systems. Finding results inconsistent with the predictions would suggest a re-evaluation of the claims made for performance in the weather prediction task.

Overview of Experiments

The experiments used the weather prediction task (WP) (Knowlton et al., 1994) employing both the observational and feedback versions of the task. In both versions of the task participants are shown between one and three discrete cues (easily discriminable tarot cards) on a trial-by-trial basis. In the feedback version participants are asked to predict a binary outcome (rainy or fine weather) and receive corrective feedback as to the actual outcome. In the observation version the outcome appears simultaneously with the cues and no overt prediction is required.

In both experiments participants completed 102 training trials either under the observation or the feedback arrangement and then completed a further 42 test trials. The test trials required predictions but no corrective feedback was provided. These test trials served two purposes: 1) they allowed us to measure learning in the observation groups (note that in the observation group no overt prediction is required during training so no training phase data are recorded), 2) they allowed us to compare performance on the same sized sample of trials after equal of amounts training in the two versions of the task. This direct comparison has not been possible in previous research because of the failure to equate the number and type of test trials (e.g. Shohamy et al. 2004a)¹.

Experiment 1 compared performance in both versions of the WP task under either concurrent memory load or no concurrent memory load conditions. Experiment 2 did not employ a concurrent task but focused solely on the two versions of the WP task and sought more sensitive measures of both self and task insight.

Experiment 1

Participants

Forty-eight undergraduate students from the University of New South Wales participated in the experiment in return for course credit. There were 30 females. The average age was 21.1 (range 17-54, $SD = 6.5$). Participants were randomly assigned to one of the four conditions.

Design and Materials

The experiment was a 2 (Training Task Type: Observation or Feedback) x 2 (Memory Load: Concurrent Task or Control) between subjects design. In both the Observation and Feedback versions of the WP task the stimuli presented to participants were drawn from a set of four cards, each with a different geometric pattern (squares, diamonds, circles, triangles). During training participants saw a total of 102 trials, on each of which they were presented with a pattern of one, two or three cards. Each trial was associated with one of two outcomes (Rain or Fine), and overall these two outcomes occurred equally often. The pattern frequencies are shown in Table 1, along with the probability of the outcome for each of these 14 patterns. The learning set was constructed so that each card was associated with the outcome with a different independent probability. For example, the probability of rain was 0.2 given the presence of the squares card (card 1), 0.4 given the presence of the diamonds card (card 2), 0.6 given the presence of the circles card (card 3) and 0.8 given the presence of the triangles card (card 4)².

In short, two cards were predictive of rainy weather, one strongly (card 4), one weakly (card 3), and two cards were predictive of fine weather, one strongly (card 1), one

weakly (card 2). Overall participants experienced identical pattern frequencies (order randomized for each participant), but the actual outcome for each pattern was determined probabilistically (so experienced outcomes could differ slightly across participants). The positions of the cards on the screen were held constant within participants, but counterbalanced across participants.

Procedure

Feedback version (Training Phase): Participants were told that they were to play a ‘game’ in which they pretended to be a weather forecaster. They were told that on each trial their task was to decide if the combination of cards presented predicted rainy weather or fine weather. Reassurance was given that although they would have to guess initially, they would be able to improve over trials. Once they had made their prediction, participants received immediate feedback as to the actual weather on that trial, and whether they were correct or incorrect (Screenshots of the experimental environment can be found in Lagnado et al., 2005).

Observation version (Training Phase): Similar instructions were given to participants in the observation condition but rather than emphasizing predictive accuracy, participants were simply told to observe and to try to learn which combinations of cards were associated with rainy or fine weather. In both the observation and feedback versions participants were told that a test phase would follow training.

Memory Load Task: Two groups performed the WP task in either the Observation or Feedback arrangement as described above. An additional two groups performed the WP task (Observation or Feedback) and a concurrent memory load task. The concurrent task was adapted from the “numerical Stroop task” used by Waldron and Ashby (2001). The task

required participants to remember which of two numbers was numerically larger and which was physically larger. Presentation schedules were arranged such that on 50% of trials both the size and value dimensions were congruent and on 50% they were incongruent. The numbers appeared simultaneously on either side of the card combinations and remained on screen for 1 second. Physically large numbers appeared in 24-point font, small in 14-point font.

In the Feedback version participants made a prediction about the weather and then received outcome feedback about the actual weather. A button then appeared prompting participants to answer a question about the number task. On clicking this button the cards and outcome disappeared from the screen and were replaced with a question about either the size or the value of the previously presented numbers. Participants answered this question by indicating whether they thought the number that had appeared on the left or the right of the cards was the numerically/physically larger one. Finally they received outcome feedback about their concurrent task response. In the Observation version the numbers, card combination and actual weather outcome all appeared simultaneously on the screen. The numbers remained for 1 second, and two seconds after their disappearance the prompt button for the number task appeared. Clicking this button removed the cards and outcome from the screen and replaced them with the size/value question about the numbers. With this design, it was ensured that participants had to hold information about the location, size and value of the numbers in memory either whilst making a prediction (feedback version) or whilst observing the card combinations and outcome (observation version).

Participants in the concurrent task conditions were given the explicit instruction that they were to try to perform the number task *perfectly* and with “what was left over” learn the weather prediction task (c.f., Waldron & Ashby, 2001). This instruction was crucial because

we wanted to be sure that the number task was the primary focus of working memory in both conditions.

Test Phase (both versions): The test phase comprised two parts. First all participants were asked to give probability ratings for each of the four cards. For each card they were asked for the probability of rainy vs. fine weather: ‘On the basis of this card what do you think the weather is going to be like?’ They registered their rating using a continuous slider scale ranging from ‘Definitely fine’ to ‘Definitely rainy’, with ‘As likely fine as rainy’ as the midpoint. After making these ratings, all participants performed the same version of the WP task without the concurrent memory load task. On each trial participants saw the card combination and made a prediction. No corrective feedback was provided. Each participant saw each of the 14 card patterns three times in a different randomized order for a total of 42 trials.

Results

The results are presented in the following order. First we report analysis of the concurrent task performance. Second we report training phase data in the WP task from the two Feedback groups (Concurrent / Control). Third we report test data from all four groups for the probability ratings. Fourth we report data from all four groups from the WP test prediction trials. Finally we present analyses of learning strategies using the mathematical modelling technique of Gluck et al. (2002).

Concurrent Task Performance

Accuracy in the concurrent memory task was high with an average across the 102 trials of 85.3% ($SD = 11.3$) in the Feedback group and 92.8% ($SD = 4.1$) in the Observation

group. This difference reached significance, $F(1, 23) = 4.60, p < .05$, although it seemed to be driven primarily by one participant in the Feedback condition whose performance was more than two standard deviations below the mean for that group (59%). Furthermore, analysis of the final block of trials indicated no significant difference between the groups $F(1, 23) = 1.19, p > .25$. This pattern of results suggests that participants in both conditions were, at least by the final block, attending to the concurrent task equally. This equality of performance in the concurrent task is important for interpreting any differences in performance on the different versions of the WP task (cf. Waldron & Ashby, 2001).

Training Phase Performance

Participants in both Feedback groups improved in their ability to predict the outcome across training trials. Feedback Concurrent improved from a mean of 57.3% correct predictions across the first 25 trials to a mean of 63.2% across the final 25. Feedback Control improved from 64.1% to 70.8% over the same trials. Although the improvements were of approximately the same magnitude a linear trend test showed that the improvement was only significant for the Feedback Control group, $F(1, 23) = 4.37, p < .05$; Concurrent, $F(1, 23) = 2.17, p > .09$. Moreover, a comparison of performance collapsed across the second half of the training trials showed a significantly higher number of correct predictions in the Feedback Concurrent group ($M = 71.6\%$) than in the Feedback Control group ($M = 62.4\%$), $F(1, 23) = 4.34, p < .05$. This pattern of results indicates that the concurrent task was impacting on participants' ability to learn the weather prediction task via feedback. Note that these results should not be interpreted as indicating that the Concurrent group learned nothing about the WP task (across training they performed at significantly above the chance level of 50%, ($t(11) = 2.83, p < .05$). Rather the results suggest that Concurrent participants learned *less* about the WP task than Control participants. Levels of performance in the Control group

were consistent with that seen in previous experiments after 100 training trials (e.g., Gluck et al., 2002; Lagnado et al., 2005).

Test Phase: Probability Ratings

At the end of the 102 training trials and prior to beginning the test prediction trials all participants judged the probability of the weather given each individual card. Figure 1 displays the mean probability ratings of rain for each card. An ANOVA with Card Type (1-4) as a within-subject factor and Training Task (observation or feedback) and Memory Load (concurrent or control) as between-subjects factors indicated a highly significant linear effect of Card Type, $F(1, 44) = 81.23, p < .001$ demonstrating clear discrimination between the objective probabilities associated with each card. There were no effects of Training Task or Memory Load ($F_s < 1$). There was, however, a significant linear interaction between Card Type and Memory Load, $F(1, 44) = 10.50, p < .01$. The interaction appears to be due to the concurrent task having a greater effect on the ratings given for the two strong predictors (Cards 1 and 4) than the two weak predictors (Cards 2 and 3). Participants in the control groups tended to make more extreme probability ratings for the strong predictors than those in the concurrent group. In sum, Figure 1 shows that participants' are able to distinguish well between the different cards and that their estimates tended to approximate the actual probability values. There is no suggestion that insight into the task structure is markedly different in the observation and feedback versions.

Test Phase: Performance

Figure 2 displays the mean proportion correct predictions across the 42 test trials for the four groups. It is immediately apparent from a visual inspection of the figure that the results obtained are opposite to that predicted by the multiple system account. The

detrimental effect of the concurrent task is much larger in the feedback condition than in the observation condition. The test performance of participants who had learned the WP task through observation was relatively unaffected by whether they had performed the concurrent task during training. In contrast, feedback participants who performed the concurrent task at training fared much worse at test than their counterparts who had only learned the WP task. This pattern was confirmed by a statistical analysis. A between-subjects ANOVA on mean proportion correct indicated that there was a significant effect of Memory Load, $F(1, 44) = 7.32, p < .05$, but no effect of Training Task type, $F(1, 44) = 2.55, p > .1$. The interaction between these two variables was significant, $F(1, 44) = 4.49, p < .05$. Simple effects analysis confirmed that performing the concurrent task during training had a much larger effect on subsequent test performance on the WP for participants in the feedback conditions (Feedback, $F(1, 22) = 12.88, p < .05$; Observation, $F < 1$). Performance in all four conditions was significantly above chance level (50%) (Feedback Concurrent, $t(11) = 2.49, p < .05$).

Strategy Analyses

Gluck et al. (2002) identified three classes of strategies that participants might use to ‘solve’ the weather prediction task. The strategies differ primarily in their complexity, as indexed by the number of cues relied upon by participants in making their predictions. The three classes are as follows:

Multi-cue maximizing strategy: this is the optimal strategy for learning the task. It involves responding to each pattern with the outcome most often associated with that pattern. A participant adopting this strategy must attend to all the cues present on a given trial and by doing so can achieve 92% correct predictions (note that this is not 100% because Patterns F, I, K and M all predict ‘fine’ and ‘rain’ equally often – see Table 1).

One-cue strategy: this is a sub-optimal strategy in which participants respond on the basis of the presence or absence of a single cue, disregarding other cues. For example a participant might respond ‘rain’ whenever the triangle card (Card 4) is present and ‘fine’ when it is not present, irrespective of what other cards are present. A participant adopting this strategy could achieve 87% correct predictions. Participants responding solely on the basis of the weaker predictors (Cards 2 or 3) could achieve 63% correct predictions.

Singleton Strategy: another sub-optimal strategy in which a participant learns only the responses associated with patterns on which a single card appears and guesses on the remaining trials. Because of the preponderance of singleton trials, this strategy achieves 65% correct predictions.

An additional strategy that we considered was the *multi-cue matching strategy*. This strategy assumes that participants distribute their responses to a pattern according to the actual probabilities associated with that pattern. For example, for pattern A across the 102 training trials a participant would respond ‘fine’ 9 times and ‘rain’ once (as opposed to the 10/0 ‘fine’ responses predicted by the multi-cue maximizing strategy). Evidence from a wide-range of probabilistic tasks suggests that participants often adopt such a matching strategy (even when it is not optimal) (e.g. Shanks, Tunney & McCarthy, 2002; Tversky & Edwards, 1966; West & Stanovich, 2003). This strategy can achieve 82% correct predictions.

We used the same method as Gluck et al. (2002) to fit participants’ learning profiles to the strategies described above. The basic procedure was to calculate to the degree to which each model fit the participant’s data using a least mean squares measure, with 0.00 indicating a perfect fit (see the appendix for details).

First we examined data from the 102 training trials in the Feedback Control and Feedback Concurrent groups. All 24 participants showed performance consistent with one of the four strategies. The fits lay between 0.00 and 0.16 with only three participants (from the

concurrent group) falling above 0.10^3 . Figure 3 shows the distribution of strategies across participants in the two groups. What is immediately apparent from the figure is the dominance of the multi-match strategy in the control group. Eighty-three percent of participants made predictions consistent with the use of this strategy. In the concurrent group adherence to this strategy fell to 42%, while adoption of the simpler singleton strategy rose from 0 in the control group to 42% in the concurrent group. A chi square analysis indicated that the distribution of strategies differed significantly between the two groups, $\chi^2 = 6.7, p < .05$.

This pattern of strategy distributions is consistent with the idea that participants who only had to learn the WP task tended to integrate information from several cues into their predictions, whereas those given the WP task and the concurrent number task tended to rely on simpler ‘singleton’ strategies. The pattern is inconsistent with the idea that learning the WP task through feedback is unaffected by the addition of a concurrent task. To the contrary, there appears to be a qualitative shift in the kinds of strategies that participants use in the two conditions.

Participants in the observational conditions did not make overt predictions during training so we could only compare model fits for observational and feedback groups using data from the 42 test trials. The majority of participants across all four conditions were poorly fit by the four models described. This failure to find good fits is consistent with previous attempts to fit the strategies to participants trained under observation conditions (Shohamy et al., 2004). “Best” fits ranged from 0.04 to 0.37 with 60% of participants only being fit with a tolerance of 0.10 or higher.

The distribution of strategies for the 19 participants whose data were consistent with one of the models tended to support the analyses of the training phase. Responses that were consistent with the more complex strategies (multi-cue-max and multi-cue-match) were

found primarily among participants trained in the feedback control condition (5 participants) whilst also found in the observation control condition (3 participants) but less so in the two concurrent conditions (2 participants each). Overall 12 of the 19 participants whose responses matched any of the four models came from the control conditions (8 feedback, 4 observational) suggesting that participants who were able to concentrate solely on the WP task during training were able to use this knowledge to adopt a consistent strategy at test.

Discussion

The most important result from Experiment 1 is displayed in Figure 2. A simple prediction of the multiple system account is that placing additional load on working memory should have a greater detrimental effect on tasks that recruit the declarative system. Figure 2 shows the opposite pattern of results. The task claimed to recruit the declarative system (observation) was relatively unaffected by the addition of a memory load task during training, whereas the task claimed to recruit the procedural system and thus be ‘working memory independent’ was significantly affected.

The strategy analyses provide some indication as to why this pattern of results might have been found. When participants’ attention was taken up with memorizing the numbers in the concurrent task, it appears that the majority of them were unable to learn more than a few simple cue-outcome associations and thus relied on the simple one cue or singleton strategies. Those in the control group overwhelmingly relied on the more complex multi-cue matching strategy which relies on integrating across different cue-outcome contingencies. It is interesting to note that this strategy accounts far better for performance than the multi-cue maximizing strategy, suggesting that indeed, participants are probability matching rather than maximizing. The dominance of the multi-cue matching strategy is consistent with recent results reported by Lagnado et al. 2005 under similar conditions. However, Lagnado et al.

reported that with greater amounts of training (200 trials) there was a tendency for participants to shift towards the more optimal maximizing strategy. This gradual shift from matching to maximizing after considerable amounts of training has been observed in other simpler probability learning tasks (e.g., Shanks, Tunney & McCarthy, 2002). The absence of singleton strategies in the control group is surprising given previous research showing a dominance of this strategy under similar conditions (e.g. Gluck et al., 2002). We can only speculate that this may have something to do with the motivation of participants⁴.

The results lend weight to the argument that the different versions of the WP task can not be classified *a priori* as “procedural” and “declarative” (Ashby & Maddox, 2005; Lagnado et al., 2005). The ‘signatures’ associated with performance mediated by these systems is conspicuously absent in these data. The probability ratings displayed in Figure 1 also provide little support for the idea that observation training results in more accurate explicit knowledge of the task structure. However, it could be argued that these retrospective global ratings are a rather crude and insensitive measure of insight (e.g., Lagnado et al., 2005; Lovibond & Shanks, 2002); furthermore they reveal nothing about participants’ insight into their own judgment processes. To tackle the problem of retrospective assessment, Experiment 2 employed more sensitive measures of insight in which multiple assessments of both task knowledge and self-insight were taken – either trial-by-trial or blocked – throughout the course of the experiment.

Experiment 2

The claim that the observation version of the WP task recruits a declarative system leads to the prediction that participants trained under observation conditions should exhibit better insight into both the structure of the task and their own judgment processes than participants who learn the procedurally mediated feedback version (Gluck et al., 2002; Shohamy et al.,

2004a). Learning in the procedural system is hypothesized to be unconscious. Experiment 2 aimed to test this simple prediction by comparing performance of a group given the observation version with one given the feedback version.

Participants

Thirty-two undergraduates from the University of New South Wales participated in the experiment in return for course credit. There were 22 females. The average age of participants was 20.1 (range 18 – 27, $SD = 2.1$).

Design & Procedure

The experiment was a two groups design with one group learning the Observation version of the WP task and the other the Feedback version. To assess explicit knowledge of task structure and self-insight, trial-by-trial ratings were elicited from participants. In the Feedback group, participants made a weather prediction and were then asked “How much did you rely on each card in making your prediction?” before receiving corrective feedback on the outcome. Ratings were made on a ‘drop-down’ menu by selecting between four options – “Greatly”, “Moderately”, “Slightly”, “Not at all”. Once participants had made their rating they were shown the actual outcome (i.e., RAINY or FINE), and then proceeded to the next trial. Previous experiments using these on-line ratings of insight have shown that making ratings does not affect the strategies that participants adopt in solving the feedback version of the task (Lagnado et al. 2005).

Participants in the Observation group did not make predictions so instead of rating reliance for making predictions they were asked, “How important is each card in determining the outcome?” and used the same 4-point rating scale to make a response. In addition to these on-line ratings, participants in both groups made probability ratings for each of the four cards.

Rather than one retrospective assessment (as in Experiment 1), probability ratings were made after 51 and 102 trials. Following the second set of probability ratings all participants completed 42 test trials, in which they made predictions but were not provided with feedback.

Results and Discussion

Learning Performance

As in Experiment 1, participants in the feedback group improved in their ability to predict the outcome. The mean proportion of correct predictions improved from 66% across the first 25 trials to 74.3% across the final 25. This improvement was confirmed by a significant linear trend, $F(1, 15) = 9.87, p < .05$.

Test Performance

Training in the observation or the feedback version made no difference to performance on the test trials. The mean proportion of correct predictions was 73.3% (SD = 7.0) following feedback training and 73.2% (SD = 6.9) following observation training. This pattern of identical performance is consistent with previous studies comparing observation and feedback versions of the WP task (Poldrack et al., 2001; Shohamy et al., 2004). The improvement in the observation group relative to the observation control group in Experiment 1 may be attributable to the addition of the trial-by-trial insight ratings. Encouraging participants to reflect explicitly on cue-outcome relations on each trial may have increased engagement and hence learning of the task.

Probability Ratings

Figure 4 displays the mean probability ratings of rain for each card. The upper panel shows the ratings made after 51 trials, the lower the ratings made after 102 trials. An ANOVA with Card Type (1-4) and Block (Trials 1-51, 52-102) as within-subject factors, and Group (Observation or Feedback) as a between-subject factor indicated a highly significant linear effect of card type, $F(1, 30) = 242.3, p < .001$ demonstrating that participants were clearly able to discriminate between the objective probability associated with each card. The main effects of group and block were not significant, (Block: $F < 1$; Group: $F(1, 30) = 1.52$). However, the two 2-way interactions between Card and Block and Card and Group both approached significance (Card x Block: $F(3, 28) = 2.77, p = .06$; Card x Group $F(3, 28) = 2.59, p = .07$). The Card x Block interaction seems to be due to the increased divergence between the ratings of the weaker predictors (Cards 2 and 3) between Blocks 1 and 2 while the ratings for the stronger predictors (Cards 1 and 4) show less of a change (or a change in the opposite direction). The Card x Group interaction seems to be due to the greater divergence between weaker predictors in the Feedback than in the Observation group, but a smaller divergence between weak and strong predictors in the Feedback than in the Observation group.

A simple account of what might be happening is that Observers are unable to discriminate between weak predictors after the first 51 trials (giving both a rating of around 50%) and over-estimate the predictiveness of the strong predictors (93% and 5% respectively for objective probabilities of 80% and 20%). By the end of the second block they have learned to discriminate the weak predictors but still show a tendency towards over-estimation of the strong predictors. Feedback participants are remarkably accurate with all four cards after the first 51 trials (showing a slight tendency for over-estimation), and this accuracy is largely maintained or improved upon in the second block with the exception of the marked overestimation of Card 3.

Trial-by-Trial Ratings

During each trial of the training phase, participants were required to rate either their reliance on each card for making their prediction (Feedback group), or the importance for each card in determining the outcome (Observation group). Following Lagnado et al. (2005) we analysed only those trials on which more than one card was present and collapsed across the ratings given for ‘strong’ (Cards 1 and 4) and ‘weak’ (Cards 2 and 3) predictors. For the analysis the training trials were collapsed into 20 blocks of 5 trials. Figure 5 plots the mean ratings for strong and weak cards collapsed across these blocks of trials.

A Card Type x Block x Group ANOVA showed a significant main effect of card-type, $F(1, 30) = 70.6, p < .001$ indicating participants’ clear ability to distinguish between weak and strong predictors in their on-line ratings. The main effect of group was also significant, $F(1, 30) = 5.66, p < .05$ indicating that ratings given by Feedback participants were higher on average for both weak and strong cards than those given by Observers. The linear trend across blocks was marginally significant $F(1, 30) = 3.93, p = .057$, suggesting a slight increase in the divergence of ratings for weak and strong predictors as training progressed. The interaction between card-type and group was significant $F(1, 30) = 9.72, p < .01$. Interestingly this seems to be due to the greater divergence between ‘strong’ and ‘weak’ cues in the ratings given by the Observation group than those given by the Feedback group (heavy lines are further apart than the softer lines). Perhaps this is because observers are surer about the *structure* of the task (from a very early stage) than feedback participants are about *reliance* on particular cues.

Supporting this idea, we found that Observers showed significant divergence between the ratings given to the Strong and Weak predictors earlier than Feedback participants. Using a criterion that required the mean importance ratings for strong and weak cards to be

significantly different (at $\alpha = .05$), for two consecutive blocks, the Observation group showed divergence after 6 blocks (30 trials), while the Feedback group reached divergence after 9 blocks (45 trials). The data appear to suggest that there is a slightly different time-course for development of insight into task structure and insight into cue-reliance.

The data clearly show that in neither condition does performance seem to be mediated by an implicit mechanism. Furthermore, although the blocked probability judgments and the on-line ratings might appear to contradict each other (Observers being less accurate in the former and more accurate in the latter) this is not necessarily the case. Note that for the blocked ratings, each card was assessed individually and a specific probability judgment was required, whereas for the on-line ratings we collapsed across the two strong and two weak cues and a single importance rating (1-4) was elicited. The blocked judgments indicate that Observers were overestimating strong cards and were unsure about weak cards – this ‘gross-level’ task knowledge would translate to the clear divergence seen in the on-line ratings. Feedback participants, in contrast, were better able than Observers to discriminate weak cards after the first 50 trials – showing more accurate insight into ‘specific-level’ task knowledge. Thus to the extent that the conditions differed at all, it appears that, contrary to the predictions of the multiple system account, insight is marginally better in the Feedback condition.

Strategy Analyses

Some researchers have argued that even when performance (and perhaps insight) are equated in the observation and feedback versions of the task, participants in the two conditions may be relying on qualitatively different strategies (e.g. Poldrack et al., 2001). To investigate this possibility we examined participants’ strategies using the same modelling technique employed in Experiment 1. Figure 6 shows the distribution of strategies for the training phase of the feedback condition. All participants were fit within a tolerance of 0.10

(range: 0.00 to 0.07) and there is clear evidence that the multi-cue match strategy dominates. As in Experiment 1, over 80% of participants given feedback training responded in a manner consistent with the use of the multi-cue match strategy. Figure 7 shows the distribution of strategies for the 42 test trials following observation and feedback training. In contrast to Experiment 1, only 18% of participants failed to be fit within a tolerance of 0.10 (range: 0.00 – 0.20). The more complex strategies again clearly dominate with approximately equal assignment to the two versions of the multi-cue strategy in the feedback condition, but a preponderance of multi-cue max in the observation condition. There is little suggestion that qualitatively different strategies are employed in the feedback and observation conditions. This interpretation was supported by a chi square analysis which indicated that the distribution of strategies did not differ significantly between the two groups, $\chi^2 = .60, p > .5$.

General Discussion

We examined a task that has been used extensively in cognitive neuroscience to demonstrate the operation of two apparently distinct learning systems. In two experiments we obtained results that were opposite to those predicted by this dual systems view. In Experiment 1 the introduction of a concurrent memory load had a greater detrimental effect on the procedural/implicit version of the weather prediction task than on the declarative/explicit version; in Experiment 2 we found highly accurate task knowledge and insight and, if anything, slightly better knowledge in the ‘implicit’ than the ‘explicit’ version.

Furthermore, strategy analyses revealed in Experiment 1 that the addition of a concurrent memory load reduced the number of participants relying on complex multi-cue strategies in the implicit version of the task. This result is predicted by the operation of a working memory dependent explicit system but not by a working memory independent implicit system. Finally, strategy analyses of the data from Experiment 2 provided no

evidence to suggest that qualitatively different strategies were used to solve the two versions of the task.

These results add weight to recent suggestions that the feedback version of the weather prediction task should not, *a priori*, be thought of as one that recruits an implicit procedural learning system (Ashby & Maddox, 2005; Lagnado et al., 2005). Following Lagnado et al. (2005) (see also Lovibond & Shanks, 2002) we suggest that the most parsimonious explanation of these results is that performance in both versions of the task is mediated by a single explicit learning mechanism. According to this account, a single declarative learning process drives participant's behavioural responses (i.e., their on-line predictions), their explicit judgments about the task structure (i.e., their blocked probability ratings), and their explicit judgments about their own cue usage or cue importance (i.e., their on-line cue ratings). This single mechanism account contrasts with the dual system account which posits that a procedural or implicit system mediates on-line predictions and a declarative or explicit mechanism mediates judgments about task structure and cue usage (Lovibond & Shanks, 2002).

A single mechanism account predicts that when participants' attention is drawn to a concurrent task (as in Experiment 1) fewer resources will be available for learning the cue-outcome contingencies in the weather prediction task. As a consequence, predictions should be less accurate and participants should be less able to adopt complex strategies. Both of these predictions are clearly supported by the results of Experiment 1. A single mechanism account also predicts the pattern of data found in Experiment 2. If, as we propose, both the observation and feedback versions are mediated by the same explicit mechanism, then one would expect the similar levels of performance, insight into task structure and cue reliance, and comparable adoption of strategies that we observed across the two versions of the task.

The results of the current experiments lend further support to the findings recently reported by Lagnado et al. (2005). In addition to finding accurate explicit knowledge of task structure and cue reliance, Lagnado et al. demonstrated a strong correlation between participants' reported and actual cue use in the feedback version of the task. By computing 'rolling regressions' across a 200 trial training phase, Lagnado et al. were able to show that the cues participants identified as influencing their judgment on a given trial corresponded with how they actually weighted the cues in their decisions (as revealed by the regression analyses). This close correspondence between rated use and actual use of cues seems most parsimoniously explained by the operation a single mechanism underlying both elements of performance.

Relation to neuropsychological research

Squire (2004) has criticized single system accounts of the kind proposed above as "making insufficient contact with biology" (p. 175). That is, such accounts fail to explain the neuropsychological dissociations that have been observed. The weather prediction task – taken as a paradigm 'implicit task' - has been highly instrumental in providing evidence for such proposed dissociations. If our claim that in fact the task is approached in an explicit manner is to be supported, then it is important to consider how we can address this neuropsychological evidence.

A single explicit mechanism account predicts that patients who have impairments to the explicit system – as amnesic individuals do – should show decrements in performance relative to matched controls. In line with this prediction Hopkins et al. (2004) reported that amnesic subjects with bilateral damage to the hippocampus (as a result of hypoxia) performed significantly worse than matched controls both early and late in training on the feedback version of the task. This result qualified an earlier study (Knowlton et al., 1994)

which reported that amnesic individuals were unimpaired in learning the preliminary stages of the task (50 first trials) but then exhibited impairments in the latter stages.

One key neuropsychological result that does appear to support the two system interpretation is the finding that Parkinson's patients are impaired on the feedback version but not on the observation version (Shohamy et al., 2004a). Such a dissociation is predicted by the two system account because Parkinson's patients suffer from disrupted basal ganglia function (the brain area thought to underlie the procedural learning system). However, recent data question this dissociation. Quallo (2005) noted a number of methodological problems with the Shohamy et al. study that could have accounted for the dissociation between performance on the observation and feedback versions (see footnote 1). To remedy these perceived shortcomings, Quallo adopted a design similar to that of Experiment 2 – in which the number and type of test trials following observation or feedback training were equated – but manipulated training regime within-subjects. Quallo reasoned that if the claimed dissociation between the types of task could be observed in the same individual it would provide much stronger evidence for the operation of separate learning systems. The result of interest was that Parkinson's patients performed equally well on *both* versions of the task. This pattern was unaffected by the order in which the tasks were performed, and furthermore Parkinson's patients were unimpaired relative to controls on the feedback version of the task.

This latter result is consistent with data reported by Moody, Bookheimer, Vanek and Knowlton (2004) in which Parkinson's patients were able to learn the feedback version of the task successfully, and were unimpaired relative to matched controls. Further data relevant to the question of whether Parkinson's patients can learn the feedback version of the task comes from a study by Shohamy et al. (2004b), in which they tested controls and Parkinson's patients over a three day period (600 trials per day). Relative to the controls the Parkinson's patients were impaired on the task, although their learning was significantly above chance,

and improved gradually throughout the three days. Their overall performance at the end of day 3 closely resembled the overall performance of controls at the end of day 1. Furthermore, the strategy model fits for Parkinson's patients on day 3 were almost identical to the model fits for the controls on day 1. This pattern of a lag in overall performance and in the strategy adopted to solve the task is perhaps most parsimoniously explained by a general deficit in the time taken to learn each cue-outcome contingency rather than a specific inability to engage a procedural learning system (Lagnado et al., 2005).

The picture that emerges from these neuropsychological studies is that, at least at the behavioural level, the case for a dual systems account may not be as strong as it appears. The claim that the observation and feedback versions recruit qualitatively different systems or that there is competition between different systems when performing the tasks does not appear to be supported by the most recent studies. In many cases it may be that dissociations in performance can be more easily accounted for by a generalized learning decrement than by disruptions to qualitatively distinct systems (see Kinder & Shanks, 2003, for further development of this argument).

Is the weather prediction task an information integration task?

The present findings and their interpretation appear to be inconsistent with work by Ashby, Maddox and colleagues which emphasizes the operation of two distinct implicit and explicit systems in category learning (e.g. Ashby & Maddox, 2005; Maddox & Ashby, 2004). Although the tasks used by these researchers are very different to the weather prediction task, the conclusions drawn often follow similar lines, and so it is informative to compare the findings from the different types of tasks (see for example Ashby & Maddox, 2005; Ashby & O'Brien, 2005).

Ashby and colleagues distinguish between rule-based tasks and information integration tasks. The former can be ‘solved’ by use of an explicit reasoning process, such as the adoption of a simple verbalizable rule of the form “Respond A if the stimulus has value X on dimension Y and B if it does not”. The latter are those tasks in which “accuracy is maximized only if information from two or more stimulus components (or dimensions) is integrated at some pre-decisional stage” (Ashby & Maddox, 2005, p.153). The optimal strategies for solving such tasks are often very difficult or impossible to describe verbally, and are thus thought to be learned via an implicit, procedural system.

As the definitions of the two tasks might suggest, investigations of rule-based and information-integration tasks have been primarily concerned with perceptual categorization (e.g., deciding if a line of particular orientation and width is a member of category A or B). A variety of studies employing ingenious manipulations have been used to attempt to dissociate the processes involved in learning the two types of tasks. The studies have tested the predictions of Ashby et al.’s (1998) COVIS model (COmpetition between Verbal and Implicit Systems) and found good evidence to support the key claims of the model. It is beyond the scope of this article to discuss the details of this model, but it is worth briefly considering the predictions it makes with regard to the manipulations we used in the present experiments.

The weather prediction task is technically an information-integration task because the optimal strategy requires information integration across a number of cues (Ashby & Maddox, 2005; Ashby & O’Brien, 2005). Thus ‘technically’ the task should be learned via the procedural implicit system. However, as the strategy analysis highlights, there are a number of other simpler strategies which still afford very good performance in the task. These simpler strategies are often easy to describe verbally (e.g., one-cue) and thus could be learned via an explicit reasoning process.

According to COVIS the procedural system is only recruited when ‘surprising’ rewards are received via trial-by-trial feedback (see Ashby et al., 1998, for details of why this is the case). Thus participants trained under observation conditions in the weather prediction task should *not* be able to recruit the procedural system and therefore should not be able to adopt a complex (optimal) multi-cue strategy.

A second prediction of COVIS is that any participant classified as using a multi-cue strategy should have no, or poor insight into the structure of the task and the strategies underlying their performance. This is because participants using the multi-cue strategy are assumed to have learned this optimal way of responding via the implicit procedural system.

Inspection of Figure 7 shows that the first prediction is not supported by our data. Over 80% of participants trained in the observation version of the task were classified as using complex multi-cue strategies at test. The second prediction regarding insight is not supported by our data either. The clear differentiation between strong and weak predictors shown by observers in Figure 5 and the probability ratings for each card shown in Figure 4 suggest that observers did indeed have a good explicit understanding of the task – even though they seemed to be using a strategy that, according to COVIS, can only be learned via an implicit procedural system.

It is always dangerous to compare the predictions of a model designed to interpret performance in very different tasks (i.e., low level perceptual categorisation tasks) with data from a ‘higher level’ probabilistic task which uses highly discriminable cues and cue combinations (i.e., the weather prediction task). There may be reasons why the particular type of ‘information-integration’ required in the WP task is not the same as that required in the perceptual tasks used by Ashby and colleagues (see Shohamy, et al., 2004b for a discussion of this issue).

Nevertheless, our findings that the majority of participants learn the weather prediction task in a manner that is consistent (according to COVIS) with the operation of a procedural system, that they do so under conditions where such a system “cannot” be engaged (i.e., observation); and exhibit good explicit knowledge of strategy and task structure, lead us to question the generality of the multiple-systems view advocated by Ashby and colleagues.

We are not making any claims about the evidence for the dissociation between rule-based and information-integrations tasks found with perceptual stimuli. Rather we are questioning to what extent these dissociations are driven by the idiosyncrasies of the specific tasks employed than by any domain general characteristics of distinct learning systems. It would be preferable to show similar patterns of dissociations across a range of tasks that satisfy the definition of rule-based and information-integration tasks rather than just those concerning low-level perceptual processes. Without such generality it is not clear whether Ashby and colleagues have tapped into two important facets of cognition or have simply identified processes involved in categorising a highly restricted set of stimuli.

Ashby and Maddox (2005) suggest that real-world information integration tasks “are common” but only provide classification of X-rays by experts as an example. It is not clear how such a task maps on to the definition of information integration tasks provided above. While it may be the case that experts have difficulty describing the information they use to decide if an X-ray shows a tumour or not, it does not necessarily follow that the information was *acquired* in an implicit procedural fashion. Rather, the novice radiologist may have learned the information in an extremely explicit, intentional manner, but the *product* of this explicit learning process may no longer be available to verbal report.

Conclusions

We scrutinized a paradigm implicit task that has been used extensively in cognitive neuroscience to support the claim for dissociable learning systems. In two experiments we found little evidence of the expected ‘signatures’ of implicit performance, and suggest that performance in both the observation and feedback versions of the weather prediction task is most parsimoniously explained by the operation of a single, explicit learning mechanism.

References

- Ashby, F.G., & Maddox, W.T. (2005) Human Category Learning. *Annual Review of Psychology*, 56,149-178.
- Ashby, F. G., & O'Brien, J. B. (2005). Category learning and multiple memory systems. *Trends in Cognitive Sciences*, 9, 83-89.
- Ashby, F. G., Alfonso-Reese, L. A., Turken, A. U., & Waldron, E. M. (1998). A neuropsychological theory of multiple systems in category learning. *Psychological Review*, 105, 442–481.
- Ashby, F. G., Ell, S. W., & Waldron, E. M. (2003). Procedural learning in perceptual categorization. *Memory & Cognition*, 31, 1114–1125.
- Ashby, F. G., Maddox, W. T., & Bohil, C. J. (2002). Observational versus feedback training in rule-based and information-integration category learning. *Memory & Cognition*, 30, 666–667.
- Eldridge, L.L., Masterman, D., & Knowlton, B.J. (2002). Intact implicit habit learning in Alzheimer's disease. *Behavioral Neuroscience*, 116, 722-26.
- Filoteo, J. V., Maddox, W. T., Salmon, D. P., Song, D. D.. (2005). Information-Integration Category Learning in Patients With Striatal Dysfunction. *Neuropsychology*, 19, 212-222.
- Gabrieli, J. D. (1998). Cognitive neuroscience of human memory. *Annual Review of Psychology*, 49, 87-115.
- Gluck, M., 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.
- Hertwig, R. & Ortmann, A. (2001). Experimental practices in economics: Methodological challenges for psychologists? *Behavioural and Brain Sciences*, 24, 383 – 451.

- Hopkins, R. O., Myers, C. E., Shohamy, D., Grossman, S., & Gluck, M. (2004). Impaired probabilistic category learning in hypoxic subjects with hippocampal damage. *Neuropsychologia*, 42, 524–535.
- Keri, S., Kelemen, O., Szekeres, G., Bagoczky, N., Erdelyi, R., Antal, A., Benedek, G, & Janka, Z (2000) Schizophrenics know more than they can tell: Probabilistic classification learning in schizophrenia. *Psychological Medicine*, 30, 149-155.
- Knowlton, B., Squire, L. and Gluck, M. (1994). Probabilistic classification leaning in amnesia. *Learning & Memory*, 1, 106 – 120.
- Lagnado, D.A., Newell, B.R., Kahan, S., & Shanks, D.R. (2005). Insight and strategy in multiple cue learning. *Manuscript submitted to Journal of Experimental Psychology: General*.
- Lovibond, P. F., & Shanks, D. R. (2002). The role of awareness in Pavlovian conditioning: Empirical evidence and theoretical implications. *Journal of Experimental Psychology: Animal Behavior Processes*, 28, 3–26.
- Maddox, W. T., & Ashby, F. G. (2004). Dissociating explicit and procedural-learning based systems of perceptual category learning. *Behavioral Processes*, 66, 309–332.
- 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, 438-42.
- Poldrack, R. A., Clark, J., Pare-Blagoev, E. J., Shohamy, D., Creso Moyano, J., Myers, C., & Gluck, M. A. (2001). Interactive memory systems in the human brain. *Nature*, 414, 546–550.
- Quallo, M. (2005). *The effect of corrective feedback on non-motor implicit learning in Parkinson's disease*. Unpublished Master's dissertation, Institute of Neurology, UCL.
- Reber, P. F., & Squire, L. R. (1999). Intact learning of artificial grammars and intact category learning by patients with Parkinson's disease. *Behavioral Neuroscience*, 113, 235–242.

- Schultz, W., Dayan, P., Montague, P. R. (1997). A neural substrate of prediction and reward. *Science*, 275, 1593-1599.
- Shanks, D. R., Tunney, R. J., & McCarthy, J. D. (2002). A re-examination of probability matching and rational choice. *Journal of Behavioral Decision Making*, 15, 233-250.
- Shohamy, D., Myers, C. E., Onlaor, S., & Gluck, M. A. (2004a). Role of the Basal Ganglia in category learning: How do patients with Parkinson's disease learn? *Behavioral Neuroscience*, 118, 4, 676–686.
- Shohamy, D., Myers, C. E., Grossman, S., Sage, J., Gluck, M. A. and Poldrack, R.A (2004b). Cortico-striatal contributions to feedback learning: converging data from neuroimaging and neuropsychology. *Brain*, 127, 851-859.
- Squire, L. R. (2004). Memory systems of the brain: a brief history and current perspective. *Neurobiology, Learning and Memory*, 82, 171-177.
- Tversky, A., & Edwards, W. (1966). Information versus reward in binary choice. *Journal of Experimental Psychology*, 71, 680-683.
- Waldron, E. M., & Ashby, F. G. (2001). The effects of concurrent task interference on category learning. *Psychonomic Bulletin & Review*, 8, 168–176.
- West, R. F., & Stanovich, K. E. (2003). Is probability matching smart? Associations between probabilistic choices and cognitive ability. *Memory & Cognition*, 31, 243-251.

Appendix

Strategy Analysis (adapted from Gluck et al., 2002)

We followed the strategy analysis procedure introduced by Gluck et al. (2002). We investigated four strategies: *multi-max*, which uses all four cues and selects the most probable outcome on each occasion; *multi-match*, which also uses all four cues, but selects an outcome in proportion to its probability; *one-cue*, which selects on the basis of just one cue; and *singleton*, which just learns patterns that contain individual cues.

For each individual participant we constructed ideal response sets for each of the four strategies, defined as the expected pattern of responses if the participant was reliably following that strategy. This was compared with the participant's actual responses by computing a model score:

$$\text{Score for model } M = \sum (\#rain_expected_{p,m} - \#rain_actual_p)^2 / \sum (\#present_p)^2$$

where $P = \text{pattern } A\dots N$; $\#present_p$ is the number of times pattern P appears in the trial set, $\#rain_expected_{p,m}$ is the number of rain responses expected to pattern P under model M , and $\#rain_actual_p$ is the actual number of rain responses the participant made in the trial set. The resultant score was a number between 0 and 1, with 0 indicating a perfect fit between model M and the participant's actual responses. The best fitting model (lowest score) was selected for each participant. These were computed for the 102 training trials.

Author Note

Ben R. Newell, School of Psychology University of New South Wales, David A. Lagnado and David R. Shanks, Department of Psychology, University College, London.

The support of the Australian Research Council (DP 0558181) is gratefully acknowledged. The work was part of the programme of the UK ESRC Research Centre for Economic Learning and Social Evolution. We thank Elia Vecellio and Tamara Cavenett for assistance in data collection. Correspondence concerning this article should be addressed to Ben Newell, School of Psychology, University of New South Wales, Sydney, 2052, Australia (Email: ben.newell@unsw.edu.au).

Footnotes

1. Shohamy et al. (2004a) gave observation participants a separate test phase in which there were 42 test trials (each of the 14 possible patterns repeated three times). For feedback participants there was no separate test phase, so the last 50 trials of training were treated as a 'test phase'. There are two problems with this design: i) feedback participants continued to receive feedback in their 'test phase' while observation participants did not, allowing the former to continue learning throughout the 'test'; ii) observation participants encountered all 14 patterns three times whereas feedback participants saw a random selection of patterns. Thus it is possible that feedback participants were not tested on all of the patterns (See Quallo (2005) for further discussion of these issues).
2. The pattern frequencies shown in Table 1 are slightly different from those in previous studies that have used 200 training trials (e.g. Lagnado et al. 2005). This was necessary to ensure that the independent probabilities of each card remained at 0.2, 0.4, 0.6 and 0.8 when computed across the 102 training trials used in the current experiments.
3. Gluck et al. (2002) suggest a tolerance level of 0.10 as indicative of a 'fit' with one of the models. However, other studies have simply used the lowest score to indicate a 'best fit' (Hopkins et al., 2004).
4. Lagnado et al. (2005) speculated that the inclusion of trial-by-trial performance related payment in their experiments may have had an effect on motivation and led to the increased prevalence of more complex strategies (c.f. Hertwig & Ortmann, 2001). The current experiments did not use trial-by-trial performance related payment but participants were told that the best performer would be rewarded with AUD \$15.

Table 1 The Learning Environment for Experiments 1 and 2.

Pattern	Cards present	Total	P(pattern)	P(Fine pattern)
A	0001	10	0.098	0.900
B	0010	5	0.049	0.800
C	0011	13	0.127	0.923
D	0100	5	0.049	0.200
E	0101	6	0.058	0.833
F	0110	4	0.040	0.500
G	0111	9	0.088	0.899
H	1000	10	0.098	0.100
I	1001	4	0.040	0.500
J	1010	6	0.058	0.167
K	1011	4	0.040	0.500
L	1100	13	0.127	0.077
M	1101	4	0.040	0.500
N	1110	9	0.088	0.111
Total		102	1.00	

Figure Captions

Figure 1: Experiment 1: Mean probability ratings (+SE) for rain given each card. The solid lines indicate the objective probability associated with each card.

Figure 2: Experiment 1: Mean Correct Predictions at test following training in either Feedback or Observation versions of the Weather Prediction Task. Concurrent refers to participants who performed a memory load task during training, Control to those who only performed the Weather Prediction Task.

Figure 3: Experiment 1: Learning Strategy allocation for participants trained in the Feedback task under either Concurrent task or Control (no concurrent task) conditions.

Figure 4: Experiment 2: Mean probability ratings (+SE) for rain given each card. Upper panel judgments after 51 trials, lower panel judgments after 102 trials. The solid lines indicate the objective probability associated with each card.

Figure 5: Experiment 2: Mean trial-by-trial cue importance (Observation) and cue reliance (Feedback) ratings averaged across blocks of five trials. The heavy lines are ratings made by Observation participants, the softer ones those made by Feedback participants. The dashed lines indicate ratings collapsed across the two weaker predictors, the solid lines ratings collapsed across the two stronger predictors.

Figure 6: Experiment 2: Learning Strategy allocation during Feedback training trials.

Figure 7: Experiment 2: Learning Strategy allocation during test trials following either Observation or Feedback training.

Figure 1




Figure 2




Figure 3




Figure 4




Figure 5




Figure 6




Figure 7

