

Running head: INSIGHT AND STRATEGY IN MULTICUE LEARNING

Insight and Strategy in Multiple Cue Learning

David A. Lagnado

University College London

Ben R. Newell

University of New South Wales

Steven Kahan and David R. Shanks

University College London

Address for correspondence:

David A. Lagnado

Department of Psychology

University College London

Gower Street

London WC1E 6BT, UK

[d.lagnado@ucl.ac.uk](mailto:d.lagnado@ucl.ac.uk)

Telephone: +44 (0) 20 7679 5389

Fax: +44 (0) 20 7436 4276

## Abstract

In multiple-cue learning people acquire information about cue-outcome relations and combine these into predictions or judgments. Previous studies claim that people can achieve high levels of performance without explicit knowledge of the task structure or insight into their own judgment policies. It has also been argued that people use a variety of suboptimal strategies to solve such tasks. In two experiments we re-examined these conclusions by introducing novel measures of task knowledge and self-insight, and using 'rolling regression' methods to analyze individual learning. Participants successfully learned a four-cue probabilistic environment and showed accurate knowledge of both the task structure and their own judgment processes. Learning analyses suggested that the apparent use of suboptimal strategies emerges from the incremental tracking of statistical contingencies in the environment. These findings have wide repercussions for the study of multicue learning in both normal and patient populations.

## Insight and Strategy in Multiple Cue Learning

A fundamental goal of cognition is to predict the future on the basis of past experience. In an uncertain world this involves learning about the probabilistic relations that hold between the available information and an outcome of interest, and integrating this into a singular judgment. Thus a stock-broker draws on various market indicators to predict the future price of a share, and a race-course pundit uses factors such as form and track condition to pick the likely winner of a race.

The underlying structure of this kind of prediction is captured in the multiple cue learning framework (Brunswik, 1943; Hammond, 1955), which focuses on how people learn from repeated exposure to probabilistic information. This paradigm has been applied in various areas of psychology, including human judgment (Brehmer, 1979; Doherty & Kurz, 1996; Klayman, 1988; for a review see Goldstein, 2004), learning and memory (Gluck & Bower, 1988; Knowlton, Squire & Gluck, 1994; Shanks, 1990) and neuroscience (Ashby & Ell, 2001; Poldrack et al., 2001).

A prototypical experimental task, and one that will be used in this paper, is the weather prediction task (Knowlton et al., 1994). In this task people learn to predict a binary outcome (rainy or fine weather) on the basis of four binary cues (four distinct tarot cards, see Figure 1). 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 in a task are made up of a representative sampling of possible card combinations. On each trial participants see a specific pattern of cards, predict the weather, and receive feedback as to the correct outcome. This enables them to gradually learn the cue-outcome relations, and thus improve the accuracy of their predictions.

There has been extensive research on the conditions under which people are able to master such tasks (Brehmer, 1980; Doherty & Balzer, 1988; Klayman, 1988). The main findings are that people perform well when the cues are few in number and linearly related to the outcome, when the content of the problem is meaningful, and when they receive sufficient trials and appropriate feedback. An important question that has received less attention concerns the relation between people's performance and their knowledge of what they are doing.

#### *The Apparent Dissociation of Learning and Insight*

An intriguing finding emerging from several recent studies is that even when people perform well in such tasks, they seem to lack insight into how they achieve this (Evans, Clibbens, Cattani, Harris, & Dennis, 2003; Gluck, Shohamy & Myers, 2002; Wigton, 1996; York, Doherty & Kamouri, 1987). A larger body of research has documented a similar finding in more naturalistic settings such as medical experts' decision-making (Harries, Evans & Dennis, 2000).

This apparent dissociation is illustrated in Gluck et al.'s (2002) study with the weather prediction task (hereafter WP task). They found that while participants attained high levels of predictive accuracy (well above chance), they demonstrated little explicit knowledge about what they were doing. In particular, in questionnaires administered at the end of the task they gave inaccurate estimates of the cue-outcome probabilities, and there was little correspondence between self-reports about how they were learning the task and their actual task performance.

This apparent dissociation between learning and insight has been taken as evidence for two separate learning systems (Ashby et al., 2003; Knowlton, Mangels, & Squire, 1996; Reber & Squire, 1999; Squire, 1994). On the one hand, an implicit (or procedural) system operates in the absence of awareness or conscious control, and is inaccessible to self-report. On the other hand, an explicit (or declarative) system requires awareness and involves analytic processing. Tasks like the WP task, which require the gradual learning and integration of probabilistic information, are generally considered to involve implicit learning. The lack of self-insight on this task is thus explained by the operation of an implicit system to which participants lack access.

It is also argued that these two systems are subserved by distinct brain regions that can be differentially impaired (Ashby et al., 2003; Knowlton et al., 1996; Poldrack et al., 2001). Thus the WP task has been used to reveal a distinctive pattern of dissociations amongst patient populations. For example,

Parkinson's disease patients with damage to the basal ganglia show impaired learning on the task, despite maintaining good explicit memory about task features (Knowlton, Mangels & Squire, 1996; Myers et al., 2003). In contrast, amnesic patients with damage to the medial temporal lobes appear to show normal learning but poor declarative memory of the task (Knowlton et al., 1996).

If correct, these conclusions have wide repercussions for everyday reasoning, and for the understanding and treatment of patients with neurological damage. However, there are several reasons to be cautious about this dual-process framework. In this paper we will focus on two fundamental issues: the measurement of insight and the analysis of individual learning.

### *The Measurement of Insight*

It is important to distinguish between someone's insight into the structure of a task (*task knowledge*) and their insight into their own judgmental processes (*self-insight*). In the case of the WP task, this translates into the difference between a learner's knowledge of the objective cue-outcome associations, and their knowledge of how they are using the cues to predict the outcome. And there is no guarantee that the two coincide. Someone might have an incorrect model of the task structure, but an accurate model of their own judgment process. Politicians seem particularly prone to this.

Though distinct notions, there is a tendency in previous research to run the two together. Thus it is not always clear whether claims about the

dissociation between insight and learning refer to a dissociation between self-insight and learning, task knowledge and learning, or both. Further, this conflation can infect the explicit tests given to participants. Questions that are designed to tap someone's insight into their own judgment processes may instead be answered in terms of their knowledge about the task. Such confusion needs to be avoided if firm conclusions are to be drawn about the relation between learning and insight.

#### *Insensitivity of Explicit Tests*

There are several other problems with the explicit tests commonly used to measure task knowledge and self-insight. First, these measures tend to be retrospective, asked after participants have completed a task involving numerous trials, and this can distort the validity of the assessments that people give. The reliance on memory, possibly averaged across many trials, can make it difficult to recall a unique judgment strategy. This is especially problematic if people's strategies have varied during the course of the task, making it hard if not impossible to summarize in one global response. In general it is better to get multiple subjective assessments as close as possible to the actual moments of judgment (cf. Ericsson & Simon, 1980; Harries & Harvey, 2000; Lovibond & Shanks, 2002).

Second, explicit tests often require verbalization, but this can also misrepresent the knowledge someone has about the task. In particular, it can lead to an under-estimation of insight, because participants may know what

they are doing but be unable to put this into words. This is particularly likely in probabilistic tasks, where natural language may not be well adapted to the nuances of probabilistic inference.

A third problem with previous tests of explicit knowledge, especially in the neuropsychological literature, is that they are too vague (Lovibond & Shanks, 2002). Rather than focus on specific features of the task necessary for its solution they include general questions that are tangential to solving the task (e.g., questions about location of cards on screen). Once again this reduces the sensitivity of the test to measure people's relevant knowledge or insight. This problem can lead to an over-estimation of insight (because someone may be able to recall features of the task that are irrelevant to good performance on it) or to under-estimation (because the questions fail to ask about the critical information).

The studies in this paper seek to improve the sensitivity of explicit tests on all these counts. We will use separate explicit tests for task knowledge and self-insight, and both will involve specific questions of direct relevance to the task. In the case of task knowledge these will concern probability ratings about cue-outcome relations, in the case of self-insight subjective ratings about cue usage. Such ratings-based tests will avoid any problem of verbalization. To tackle the problem of retrospective assessments we will take multiple judgments during the course of the task (either blocked or trial-by-trial).

*Analyses of Judgment and Learning*

Claims about the dissociation between insight and learning also depend on an appropriate analysis of learning performance. In tasks with a dichotomous outcome (such as the WP) participants predict the outcome on a trial-by-trial basis, and learning performance is measured in terms of correct predictions. Standard analyses then average across both individuals and trials to produce a mean percentage correct for the whole task. While this kind of approach is useful for broad comparisons across different tasks, it does not provide much information about the learning process. It ignores the possibility of individual differences in judgment or learning strategies, and that these strategies may evolve or change during the course of the task.

*The Lens Model*

A richer approach to the analysis of the judgment process is provided by the Lens Model framework (for overviews see Cooksey, 1996; Goldstein, 2004). This is founded on the idea that people construct internal cognitive models that reflect the probabilistic structure of their environment. A central tenet of this approach is that people's judgmental processes should be modeled at the individual level before any conclusions can be drawn by averaging across individuals. This is done by inferring an individual's judgment policy from the pattern of judgments that they make across a task. More specifically, a judge's policy is captured by computing a multiple linear regression of their judgments onto the cue values across all the trials in the

task. The resultant beta-coefficients for each cue are then interpreted as the weights that the judge has given to that cue in reaching their judgments (cue utilization weights).

Each judge's policy model can then be assessed against the actual structure of the task environment. This is done by computing a parallel multiple linear regression for the actual outcomes experienced by the judge onto the cue values (again across all task trials). The resultant beta-coefficients are interpreted as the objective cue weights for the judge's environment. If all the participants have been exposed to the same environment then the objective cue weights revealed by this computation will be the same for everyone. However, this technique allows for the possibility that different individuals experience different environmental structures.

A judge's policy (their cue utilization weights) can then be compared with the objective weights to see how well they have learned the task environment. This is illustrated by the Lens Model (see Figure 2), in which one side of the lens represents the structure of the environment, the other side represents an individual's cue utilization. The computation of these objective and subjective weights also facilitates the assessment of explicit judgments. Thus to measure task knowledge an individual's explicit ratings of the cue-outcome relations can be compared with the objective weights, and to measure self-insight their explicit ratings of their own cue usage can be compared with the cue utilization weights.

The Lens Model framework thus provides a means to analyze individual judgmental processes. However, although it avoids the loss of information incurred by averaging over participants, it still loses information by averaging over trials. It fails to capture the dynamics of a learning task – both in terms of potential changes in the environment, and potential changes in a judge’s policy. In particular, the reliance on global weights ignores the fact that both the actual weights experienced by the judge, and the judge’s own subjective weights, may vary across the course of the task.

This is a problem even when the underlying structure of the environment is stationary (as it is in the WP task), because the cue-outcome patterns that someone actually experiences (and therefore the environmental weights) may not be representative of the underlying probabilistic structure, especially early on in a task. Analyzing individual judgment policies just in terms of their averaged performance across all the trials ignores this possibility, and as a consequence may under-estimate someone’s performance. Moreover, it overlooks the possibility that the person’s judgment policy may change over trials, and that such changes may track variations in the actual environment.

A related shortcoming is that these global analyses assume that the judge has a perfect memory for all task trials, and that they treat earlier trials in the same way as later trials. But both of these assumptions are questionable

– people may base their judgments on a limited window of trials, and may place more emphasis on recent trials (Slovic & Lichtenstein, 1971).

### *Dynamic Models of Learning*

The need for dynamic models of learning is now widely recognized, and a variety of different models are being developed (Dayan, Kakade & Montagne, 2000; Friedman, Massaro, Kitzis & Cohen, 1995; Kitzis et al., 1998; Smith et al., 2004). A natural extension to the lens model is the ‘rolling regression’ technique introduced by Kelley and Friedman (2002) to model individual learning in economic forecasting. In their task participants learned to forecast the value of a continuous criterion (the price of Orange juice futures) on the basis of two continuous-valued cues (local weather hazard and foreign supply). Individual learning curves were constructed by computing a series of regressions (from forecasts to cues) across a moving window of consecutive trials. For example, for a window size of 160 trials, the first regression is computed for trials 1 to 160, the next for trials 2 to 161, and so on. This generates trial-by-trial estimates (from trial 160 onwards) for an individual’s cue utilization weights, and thus provides a dynamic profile of the individual’s learning (after trial 160).

Each individual learning profile is then compared with the profile of an ‘ideal’ learner exposed to the same trials. Regressions for each ideal learner are also computed repeatedly for a moving window of trials, but in this case the actual criterion values (prices) are regressed onto the cues. The estimates

of the ideal learner thus correspond to the best possible estimates of the objective cue weights for each window of trials.

The rolling regression technique thus provides dynamic models of both actual and ideal learners, and permits trial-by-trial comparisons between the two as the task progresses. For example, in analyzing the results in their Orange Juice task, Kelley and Friedman (2002) compared actual and ideal learning curves to show that while ideal learners converged quickly to the objective weights, participants learned these weights more slowly, and their final predictions tended to over-estimate the objective cue weights.

In this paper we will use similar techniques to analyze individual judgment policies. We will also use smaller window sizes than those used by Kelley and Friedman, in order to simulate a more realistic memory constraint. This should provide a finer-grained analysis of the dynamics of individual learning.

### *Strategy Analyses*

A somewhat different approach to modeling individual learning has been developed by Gluck et al. (2002). Based on post-experiment questionnaires they identified three main strategies for learning the WP task: (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.

Gluck et al. (2002) constructed ideal judgment profiles for each of these strategies, and then fit these to the actual judgments given by each participant. Across all 200 trials the best fitting model was the singleton strategy (fitting 80% of participants). However, when the same analysis was conducted across blocks of 50 trials it revealed a gradual shift from singleton to multi-cue strategies (by the final block singleton and multi-cue strategies each fitted around 40% of participants).

This approach marks an improvement over previous analyses of the WP task, and provides a means of modeling individual learning strategies, as well as possible shifts in these strategies during the course of the task. It also resonates with recent work on the adoption of simple heuristics (Gigerenzer et al., 1999; Kahneman, Slovic & Tversky, 1982). However, the modeling suffers from the use of global (objective) cue-outcome associations to construct the ideal judgment profiles for each strategy. This ignores the possibility that participants encounter a non-representative sampling of the environment early on in the task, and thus may under-estimate the number of participants adopting a multi-cue strategy. For example, in fitting the multi-cue strategy it is assumed that a participant knows the correct cue-outcome associations from the beginning of the task. As noted above this is an unrealistic assumption. The rolling regression method overcomes this

shortcoming, because it compares an individual's judgments with those that would be made by an ideal judge exposed to the same information (and thus possibly to a deviant sample of cases in which the experienced cue-outcome relations are quite different from the objective ones).

Another feature of the rolling regression technique is that it offers a more parsimonious explanation for Gluck et al.'s findings. The existence of several strategies, and the apparent shift in strategies as the task progresses, can be accounted for by the operation of a single generalized learning procedure. One-cue or singleton strategies correspond to cases where a single cue-outcome association dominates (usually early in learning), and the shift from these simpler strategies to multi-cue strategies emerges once the other cue-outcome weights are sufficiently learned. The existence of both strong and weak cue-outcome weights in the task will promote this pattern: the stronger cues will be learned earlier on, and the weaker cues later on (if at all).

In short, the multiple strategies identified by Gluck et al.'s analyses may emerge from a more general learning process. By using the rolling regression technique the viability of this account can be tested. It also has the benefit of supplying an appropriate ideal (normative) standard against which actual learning strategies can be assessed.

In summary, this paper aims to extend our understanding of how people learn in multiple cue tasks, and re-examine the prevalent claim that good performance can be achieved in the absence of insight. In doing so, it

will introduce more appropriate measures of both task knowledge and self-insight, and more fine-grained methods for analyzing the dynamics of individual learning.

### Overview of Experiments

Both experiments are based on the weather prediction task (Knowlton et al., 1994). This task requires participants to predict a binary outcome (rainy or fine weather) on the basis of four binary cues (the presence or absence of four tarot cards). Each card is independently associated with the outcome with a different probability, and overall each outcome occurs equally often. By learning these independent cue-outcome associations, participants can improve their predictive accuracy throughout the task. However, the probabilistic nature of the environment means that the best participants can expect to achieve is 83% correct predictions.

We have introduced new test questions into the basic WP paradigm so as to provide more precise tests of both task knowledge and self-insight. To measure task knowledge participants are asked to judge the probability of the outcome (rainy vs. fine weather) on the basis of each individual card. These judgments can then be compared with the actual probabilistic relations between cards and outcome. To measure self-insight participants are asked to rate how important each card was for making their predictions. These ratings can be compared with the objective predictive value of each card. Experiment

2 also introduced a novel on-line measure of self-insight. Details are provided in the introduction to that experiment.

## Experiment 1

### *Method*

#### *Participants and Apparatus*

Sixteen students from University College London took part in the study. They were paid a basic turn-up fee of £2, and received additional payment depending on their performance in the task. The entire experiment was run on a laptop computer using a software program written in Visual Basic 6. Participants were tested individually in a sound-proofed room.

#### *Materials*

The stimuli presented to participants were drawn from a set of four cards, each with a different geometric pattern (squares, diamonds, circles, triangles; see Figure 1). Participants saw a total of 200 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).

In short, two cards were predictive of rainy weather, one strongly (card 1), one weakly (card 2), and two cards were predictive of fine weather, one strongly (card 4), one weakly (card 3). 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 position of the cards on the screen were held constant within participants, but counterbalanced across participants.

### *Procedure*

At the start of the experiment participants were presented with the following on-screen instructions:

Thank you for agreeing to take part in this experiment. In the experiment you will be playing a “game” in which you pretend to be a weather forecaster. On each trial you will see between one and three “tarot cards” (cards with squares, diamonds, circles or triangles drawn on them). Your task is to decide if the combination of cards presented predicts RAINY weather or FINE weather. At first you will have to guess, but eventually you will become better at deciding which cards predict RAINY or FINE weather. As an incentive to learn how to predict the weather, you will be paid 1p for every correct forecast you make. A scale on the screen will display your current earnings throughout the experiment. There are 200 trials in total. After each

block of 50 trials there will be a short break during which you will be asked some questions about the task and your performance.

After reading the instructions participants moved onto the first block of 50 trials. They initiated a trial by clicking on a button labelled 'New trial'. On each trial a specific pattern of cards (selected from Table 1) was displayed side-by-side on the screen (see Figure 3). Participants were then asked to predict the weather on that trial, by clicking either on RAINY or FINE. 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. Correct responses were signalled with a thumbs-up sign, and a 1 pence increment in the on-screen earnings indicator. Incorrect responses were signalled with a thumbs-down, and no change in the earnings indicator.

At the end of each block of fifty trials participants answered two different sets of test questions. In the *probability* test 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. In the *importance* test participants were asked how much they had relied on each card in making their predictions: 'Please indicate how important this card was for making your predictions'. They registered their rating using a continuous

slider scale ranging from 'Not important at all' to 'Very important', with 'Moderately important' as the midpoint. Both tests were given at the end of each block. After the final test phase participants were informed of the total amount of money they had earned. They were then paid and debriefed.

### *Results and Discussion*

#### *Learning Performance*

Across the task participants steadily improved in their ability to predict the outcome. The mean proportions of correct predictions for each block of fifty trials are shown in Figure 4 (top panel). A linear trend test showed a significant improvement across blocks,  $F(1, 15)=10.3, p < 0.05$ , and by the final block mean performance approached the optimal level of 83% correct.

#### *Probability Ratings*

After each block of fifty trials participants judged the probability of the weather for each individual card. The mean probability ratings of Rain for each card across the four blocks are shown in Figure 5 (top panel). An ANOVA with card type (1-4) and block (1-4) as within-subject factors revealed a main effect of card type,  $F(3, 15) = 16.79, p < 0.001$ , no effect of block,  $F(3, 15) = 0.64, ns.$ , and an interaction between card type and block,  $F(9, 15) = 2.65, p < 0.01$ .

Inspection of Figure 5 shows that participants improved in their ability to discriminate between the probabilities of each card through the course of the task. Recall that the actual probabilities of Rain for each card were 0.2, 0.4,

0.6 and 0.8 for Cards 1-4 respectively. By block 4 mean probability estimates were close to the actual values, except for Card 4, where the actual value was slightly over-estimated.

The observed interaction between card type and block suggests that participants learn about strong cards (and discriminate between them) sooner than they do for weak cards. This is supported by the fact that ratings for card 1 and card 4 differ significantly by block 1,  $t(15) = 4.18, p < 0.01$ , whereas ratings for card 2 and card 3 do not differ on block 1,  $t(15) = 0.36, ns.$ , and only differ significantly by block 3,  $t(15) = 2.39, p < 0.05$ .

In sum, throughout the course of the task participants gradually learn the probabilistic relations between cards and outcomes, and they tend to learn about the strong cards (cards 1 & 2) sooner than the weak cards (cards 2 & 3).

### *Cue Usage Ratings*

After each block participants rated how much they had relied on each card in making their predictions. The main question of interest here is whether participants' usage ratings discriminated between strongly and weakly predictive cards. Thus we combined ratings for the two strongly predictive cards (Card 1 and Card 4) into one group (*Strong*), and ratings for the two weakly predictive cards (Card 2 and Card 3) into another group (*Weak*). The mean cue usage ratings for each group across the four blocks are shown in Figure 6 (top panel). An ANOVA with card strength (weak vs. strong) and block (1-4) as within-subject factors revealed a main effect of card

strength,  $F(1, 15) = 7.11$ ,  $p < 0.05$ , no effect of block,  $F(3, 15) = 0.78$ , *ns.*, and a marginal interaction between card strength and block,  $F(3, 15) = 2.77$ ,  $p = 0.05$ .

Paired comparisons revealed that the difference between strong and weak groups was not significant for block 1,  $t(15) = 0.09$ , but was marginally significant by block 2,  $t(15) = 2.03$ ,  $p = 0.06$ , and significant for block 3,  $t(15) = 2.54$ ,  $p < 0.05$ , and block 4,  $t(15) = 3.85$ ,  $p < 0.01$ . These tests confirm that as the task progressed participants reported that they relied more on strong cards than weak ones.

#### *Rolling Regression Analyses*

A finer-grained analysis of individual learning is provided by the use of rolling regressions. This technique was introduced by Kelley and Freidman (2002) to model continuous predictions. In order to apply it to the binary predictions made in the WP task we used logistic rather than linear regression. We also used a smaller window size (50 rather than 160 trials) to provide a more realistic memory constraint and to capture the dynamical fluctuations in the environment.

Two trial-by-trial learning profiles were constructed for each participant, one that captured their actual judgment policy (*actual judge*) and the other that captured the policy of an ideal learner exposed to the same information (*ideal judge*). Both learning profiles were generated by computing a sequence of consecutive logistic regressions (from trials 51 through to 200) across a moving window of 50 trials. Thus a participant's actual profile was

made up of four beta-weight curves, one for each card, tracking the relation between the card values on each trial and the outcome that the participant predicted on that trial. The ideal profile for each participant was also made up of four beta-weight curves, in this case tracking the relation between the card values and the actual outcomes. It is termed 'ideal' because each beta-weight corresponds to the best possible (LMS) estimate for the weight given the experienced data (cf. Kelley & Friedman, 2002).

This technique allows us to compare the evolution of an individual's judgment policy with that of an ideal learner exposed to the same trials. In Figure 7 the judgment profiles for two participants are shown. For ease of presentation ideal and actual curves are shown just for the two strong cards (1 and 4). In each plot, the regression weights on the y-axis correspond to the strength with which that card predicts the outcome (with +ve values predicting Fine, and -ve values predicting Rain). The regression weight for a given card corresponds to the log-likelihood ratio for the outcome given the presence of that card.

$$Weight(card_j) = \ln \frac{P(fine | card_j)}{P(rain | card_j)}$$

This can be translated into the odds (or probability) of the outcome given a specific card. For example, a weight of +2.2 for card 4 corresponds to a likelihood ratio = 9.03, and thus to approximate odds of 9:1 in favour of Fine (probability = 0.9).

It is apparent from Figure 7 that the learning curves for each participant (actual judge) are qualitatively close to the ideal curves. The same curves were constructed for all 16 participants, and all showed a similar pattern. In order to draw some general conclusions the individual regression models for each participant can be combined to yield averaged learning curves. It should be noted that this is not the same as averaging all the learning data first and then fitting one statistical model. Also, by inspecting the individual models we can check whether the averaging process is representative of the individual profiles, and whether it obscures any individual differences.

There are two qualitative aspects of the task environment that participants need to master: the *direction* and the *strength* of the four cards. The direction of a card corresponds to the outcome that it predicts. Thus cards 1 and 2 share the same direction because they both predict Fine weather, and cards 3 and 4 share the opposite direction because they predict Rain. To assess whether participants' learning profiles are sensitive to card direction we can compare the regression curves for each card. The averaged ideal and judged curves for all 16 participants for the two strong cards are shown in Figure 8 (top panel), and for the two weak cards in Figure 9 (top panel).

Inspection of these figures shows that participants' learning curves mark the right direction for both weak and strong cards. They also show that the curves are qualitatively similar to the equivalent curves for the ideal

learners, but with a tendency to over-estimate towards the end of the task.

This general pattern is also present in the majority of the individual learning profiles.

The second main feature of the task environment is the difference in strength between the cards. As noted above, cards 1 and 4 are strongly predictive, cards 2 and 3 are weakly predictive. To establish whether participants' learning profiles distinguish between strong and weak cards we have combined their regression curves for strong vs. weak cards (see Figure 10). It is clear from Figure 10 that participants do discriminate between the strength of the cards, and these 'implicit' measures of policy correspond well with the explicit measures (the reported cue usages shown in Figure 6).

In sum, the averaged learning curves confirm that participants' judgment profiles capture the two main qualitative aspects of the environment, the predictive direction and strength of the four cards. They also show that qualitatively participants are doing as well as the ideal judges. However, both the individual and averaged profiles reveal a tendency for the actual judges to over-estimate the regression weights relative to the ideal judges. This anomaly is similar to that found by Kelley and Friedman (2002) and will be analysed in the general discussion.

### *Summary of Experiment 1*

Across a range of measures and analyses participants performed well on the WP task, and demonstrated both accurate task knowledge and good

self-insight. This contrasts with previous research that claimed a dissociation between learning performance and task insight.

By taking repeated judgments, and constructing dynamic learning models, we can also address the question of whether participants use different strategies, and shift their strategy as the task progresses. The blocked probability judgments (Figure 5) and importance ratings (Figure 6) suggest that participants learn to use the strongly predictive cards first, and gradually learn to use the weaker cards. This is also confirmed by the dynamic learning curves (Figure 10). This alternative to the multiple strategy view will be discussed in more detail in the general discussion.

## Experiment 2

In contrast to several previous studies (e.g., Gluck et al., 2002) the results in Experiment 1 showed that participants have good task knowledge and self-insight when completing the WP task. The rolling regression techniques supported these conclusions, and suggested that participant's individual judgment policies were qualitatively similar to those of an ideal judge. Experiment 2 sets out to replicate this finding while introducing a trial-by-trial measure of participant's cue usage. This involves asking participants to rate how much they rely on each card straight after they have made their prediction. A similar technique has been used by Harries and Harvey (2000), and the resultant cue weightings made by each participant were closely correlated with their implicit judgment policies (as revealed by

regression techniques). However, this pattern was only found when participants were engaged in an advice integration task, not with a contingency-based task (such as the WP task). One difference between their study and the current one is that they compared participants' self-ratings against regression models computed across the global set of trials, rather than a moving window. This ignores the fact that participants actually experience a dynamic learning environment, and may have a limited memory window. A second difference is that in their study the cues and outcome were continuous-valued rather than binary. Both of these differences may account for their finding that participants lacked insight in the contingency-based task.

### *Method*

#### *Participants and Apparatus*

Sixteen students from University College London took part in the study. Payment and testing conditions were the same as in Experiment 1.

#### *Procedure*

The stimuli and instructions were identical to Experiment 1 except for the addition of a cue rating stage after each prediction. On each trial, just after the participant had predicted the weather outcome, a drop-down menu appeared beneath each card present on that trial. Participants were asked "How much did you rely on each card in making your prediction?" and had to select between four options – "Greatly", "Moderately", "Slightly", "Not at

all". Once they had registered their judgments they were shown the actual outcome, and then continued to the next trial.

### *Results and Discussion*

#### *Learning Performance*

As in Experiment 1 participants improved in their ability to predict the outcome. The mean proportions of correct predictions for each block of fifty trials are shown in Figure 4 (bottom panel). In contrast to Experiment 1, however, there is a tail-off in performance towards the end of the task. This is confirmed by a significant quadratic trend,  $F(1, 15) = 12.17, p < 0.05$ , but no linear trend,  $F(1, 15) = 3.51, ns$ .

This tail-off in learning performance is most likely due to the extra cognitive effort introduced by the trial-by-trial cue usage ratings. Experiment 2 took longer to complete than Experiment 1, and participants may be showing signs of fatigue towards the end of the task.

#### *Probability Ratings*

The mean probability ratings of Rain for each card across the four blocks are shown in Figure 5 (bottom panel). An ANOVA with card type (1-4) and block (1-4) as within-subject factors revealed a main effect of card type,  $F(3, 15) = 30.97, p < 0.0001$ , no effect of block,  $F(3, 15) = 0.88, ns.$ , and an interaction between card type and block,  $F(9, 15) = 2.07, p < 0.05$ .

Inspection of Figure 5 shows that participants improved in their ability to discriminate between the probabilities of each card through the course of

the task. Once again their estimates tended to approximate the actual probability values, with the exception of the over-estimation of one card.

The observed interaction between card type and block suggests that strong cards are learned about (and discriminated between) sooner than weak cards. This is supported by the fact that ratings for card 1 and card 4 differ significantly by block 1,  $t(15) = 7.69, p < 0.001$ , whereas ratings for card 2 and card 3 do not differ on block 1,  $t(15) = 0.01, ns.$ , and only differ significantly by block 3,  $t(15) = 3.41, p < 0.01$ .

In sum, as with Experiment 1, participants gradually learn the probabilistic relations between cards and outcomes, but tend to learn about the strong cards (cards 1 & 2) sooner than the weak cards (cards 2 & 3).

#### *Cue Usage Ratings (Blocked)*

The mean cue usage ratings for weakly vs. strongly predictive cards across the four blocks are shown in Figure 6 (bottom panel). An ANOVA with card strength (weak vs. strong) and block (1-4) as within-subject factors revealed a main effect of card strength,  $F(1, 16) = 13.01, p < 0.01$ , no effect of block,  $F(3, 16) = 0.77, ns.$ , and an interaction between card strength and block,  $F(3, 16) = 4.83, p < 0.01$ .

Paired comparisons revealed that the difference between strong and weak groups was not significant for block 1,  $t(15) = 0.47$ , but was significant by block 2,  $t(15) = 4.69, p < 0.01$ , and for block 3,  $t(15) = 4.37, p < 0.01$ , and block 4,  $t(15) = 2.89, p < 0.05$ . This replicates the finding in Experiment 1 that as

the task progressed participants reported that they relied more on strong cards than weak ones.

#### *Cue Usage Ratings (Trial-by-Trial)*

In addition to blocked cue usage ratings participants also made similar ratings on a trial-by-trial basis. As before these were grouped into ratings for strong and weak cards. Figure 11 presents the mean ratings for strong and weak cards, where values on the y-axis represent how much participants said they relied on each card, with 4 = "Greatly", 3 = "Moderately", 2 = "Slightly", 1 = "Not at all". Inspection of this figure reveals that from early on in the task participants rated strong cards as more important than weak ones. Indeed if ratings are averaged across trials 10-20, there is a significant difference between strong and weak cards,  $t(15) = 2.22, p < 0.05$ . It should be noted, however, that because we have averaged across participants it does not follow that any particular individual makes this discrimination so early on in the task. Nevertheless this explicit measure supports the conclusion that participants develop insight into their own cue usage relatively early in the task.

#### *Regression Analyses*

The same analyses were carried out as in Experiment 1. Figure 12 shows the judgment curves for two representative participants, and the bottom panels of Figures 8 and 9 show averaged judgment curves for strong cards and weak cards respectively. Figure 10 (bottom panel) shows the

averaged judgment curves for strong vs. weak cards. Note the similarity between participants' implicit trial-by-trial cue usage for strong vs. weak cards in Figure 10 (bottom panel) and their explicit trial-by-trial judgments about their cue usage in Figure 11. This suggests a nice fit between what participants say they are doing and what they are actually doing.

Overall the judgment profiles reveal the same pattern as in Experiment 1: participants gradually learn to distinguish the direction of the cards, and discriminate between strong and weak cards. They also show the same tendency to over-estimate the regression weights relative to the ideal learner.

### General Discussion

In two experiments we investigated how individuals master a multicue probability learning task. The central findings were that participants achieved near optimal performance, and that this was accompanied by accurate task knowledge and self-insight. This contrasts with previous studies (Evans et al., 2003; Gluck et al., 2002; Wigton, 1996; York et al., 1987), which claim that people achieve high levels of performance in the absence of explicit knowledge of the task structure or their own judgment policy. The investigation of self-insight was extended in the second experiment by taking trial-by-trial reports of cue usage immediately after participants had made their predictions. Even early in the task these online reports distinguished between strong and weak cues.

To provide a more fine-grained analysis of the dynamics of learning we modelled each individual's learning using a novel 'rolling regression' technique. This generated implicit usage profiles for each cue across the course of the task, taking into account a limited memory capacity. These implicitly derived measures of learning supported the conclusions drawn from the explicit measures – participants were sensitive to the structure of the learning environment, and from a relatively early stage in the task their implicit cue weights distinguished both the directions of each cue and their relative strengths. Moreover, a tendency to overshoot the ideal cue weights towards the end of the task was mirrored by the over-estimation of the explicitly stated probabilities (for discussion see below).

The close fit between derived and explicit measures is well illustrated if one compares the regression profiles for participants' implicit cue usages (Figure 10) with their explicit trial-by-trial importance ratings (Figure 11). Both show a clear discrimination between the strong and weak cues that increases as the task progresses. This demonstrates that from early on in the task participants were accurate in their reports about their own cue usage. Again this flies in the face of received wisdom about the frailties of people's self-reports.

These findings raise serious questions for the received view that probabilistic category learning is purely an implicit task, and that people can master it the absence of awareness and insight. In particular, the standard

assumption that such learning involves an implicit procedural system inaccessible to conscious awareness is compromised by our finding that participants' explicit task knowledge corresponds closely with the actual structure of the learning environment, and that their knowledge of what cues they use corresponds well with how they actually weight the cues (as revealed by the regression analyses). More generally, our findings are problematic for the standard distinction between explicit and implicit learning systems. According to the classifications introduced and defended by proponents of this distinction, we have shown that a paradigm implicit task is being solved by participants in an explicit manner. This has strong repercussions for the various neuropsychological studies that have built upon this alleged distinction (see below).

### *Models of Self-Insight*

Given that most previous research has claimed that people lack self-insight, there are few positive proposals about its nature. Most theorizing has focused on dual-route models that posit separate mechanisms for people's verbal reports and their non-verbal behaviour. An old but still influential view is provided by Nisbett and Wilson (1977). They argue that when people report how particular cues have influenced their responses they do not access internal states, but base their reports on a priori causal models derived from their beliefs about what people typically do (or ought to do) in such situations. This view meshes well with the general implicit-explicit distinction

– the way people actually use the cues is implicit (and inaccessible to awareness) whereas their post experiment reports are explicit but incorrect. Irrespective of the merits of Nisbett and Wilson's account (for criticisms see Ericsson & Simon, 1984; White, 1988), it is not applicable to the kind of learning task studied here. The abstract nature of the stimuli makes it unlikely that participants have an a priori causal model about what people would (or should) do in such tasks.

An alternative approach is to take seriously the possibility that people have access to the internal states that influence their behaviour, and that this drives both their online predictions and their explicit verbal reports. This would explain the strong correlation between participants' implicit cue weightings (as revealed by the regression models), their explicit cue usage ratings (both online and blocked), and their blocked probability judgments. It would also account for the tendency (noted in the introduction) to conflate task knowledge and self-insight, because on this model both share the same common cause. Thus although conceptually distinct, participants' explicit judgments about the structure of the task will usually be highly correlated with their expressed cue usage. This is particularly true when they perform well, because high performance requires that their subjective cue weightings correspond to the objective cue weights in the task environment. In short, we conjecture that people gradually acquire cue-outcome weightings that veridically represent the environment to which they are exposed, and that

these weightings drive their online predictions and their explicit reports about both the task structure and their own cue usage.

We believe this to be the most parsimonious explanation of the current data. There are, however, several other models that have been proposed in the literature. Following Lovibond and Shanks (2002) we can distinguish between two major classes of model: single process models that posit just one learning mechanism, and dual process models that posit two independent learning mechanisms, one declarative and the other procedural.

Applied to our current studies, the two classes of model are depicted in Figure 13. Model A asserts a single declarative learning process that drives people's behavioural responses (e.g., their online predictions), their explicit judgments about the task structure (e.g., their probability ratings), and their explicit judgments about their own cue usage (both blocked and online). In contrast, model B asserts that people's online predictions derive from a procedural learning process, and their explicit judgments derive from a separate declarative learning system.

In a broad sense both models are consistent with the data from our current studies. However, the dual-process model is challenged by the close concordance between the explicit measures and online predictions. This model predicts that explicit measures and online predictions will sometimes dissociate, but there is no evidence of this in our studies (and little support in the wide range of studies surveyed by Lovibond and Shanks). Thus, not only

is the single process model favoured on grounds of simplicity, insofar as it posits one rather than two learning systems, but it is also supported by the close match between behavioural and explicit measures.

A more direct way to discriminate between these models would be to examine whether disruption to the declarative learning system inhibits behavioural responses. If the declarative system can be impaired or disrupted, without compromising online predictions, this would be strong evidence against a single process model. Thus far there is little evidence in support of such a possibility. For example, although Knowlton et al. (1994, 1996) claimed that amnesics with disrupted declarative systems could still master the WP task, more recent research has shown that amnesics are severely impaired on the task (Hopkins, Myers, Shohamy, Grossman & Gluck, 2004).

#### *Awareness of Combination Rule*

So far we have given an account of how people's conscious awareness of the cue-outcome contingencies influences their predictions. This leaves open the question of whether people are aware of the judgment rule used to combine this information into a single prediction. This is relatively trivial in the current task, because the final outcome is an additive combination of each individual cue, but one might ask whether people could learn to use more complex combination functions (e.g., non-linear ones), and maintain explicit knowledge of the type of rule they are using. This question is complicated by the fact that in more complex situations configural strategies might be

appropriate. Whatever the answer to such questions, the additive combination rule in the current task is representative of a vast number of real world tasks (Hastie & Dawes, 2001). Thus the conclusions reached here have a wide general applicability.

### *Strategy Analyses*

Gluck and colleagues (Gluck et al., 2002; Shohamy et al., 2004) have recently suggested that people adopt a variety of strategies to solve the weather prediction task. In particular they identify three potential strategies: a multi-cue strategy that uses all the cue-outcome relations, a single-cue strategy that uses just one cue, and a singleton strategy that only uses a cue when it is presented on its own. Despite the suboptimality of the latter two strategies, they can still lead to above chance performance. On the basis of model fits to their empirical data Gluck et al. (2002) make two claims: that the majority of participants adopt a singleton strategy, and that there is a gradual shift in strategy from singleton to multi-cue.

The empirical data and regression models from our two experiments, however, suggest a simpler account of this apparent strategy use and strategy shift. Both features would arise naturally from the operation of a general (flexible) learning mechanism that incrementally learns the cue-outcome associations for all cues. The key idea here is that as people gradually learn these relations they might look as if they initially use a singleton strategy, and only later shift to a multi-cue strategy. But this pattern of responding could

equally result from a tendency to learn about the strong cues first, and only learn about the weaker cues later in the task. This is compounded by the fact that early in the task the experienced cue-outcome contingencies will sometimes be unrepresentative of the actual contingencies, especially for the weaker cues.

This alternative explanation is confirmed both by the explicit measures of participants' knowledge and by the individual regression analyses. In both cases participants learned to distinguish between all four cues (with respect to strength and direction) relatively early, and knowingly relied more on the strong cues than the weak ones. In addition, the claim that the majority of people use a singleton strategy seems independently implausible. It is unlikely that people are able to learn each individual cue-outcome relation, but are unable to make any use of this information on trials where more than one cue is present. For example, once a participant has learned that card 1 predicts rain, and that card 2 predicts rain, it is unlikely that when faced with a pattern consisting of both card 1 and card 2 they would guess the outcome randomly rather than predict rain. Indeed our learning data show that from early on in the task participants are predicting more successfully than the maximum possible on a singleton strategy.

In sum, the apparent strategy usage and shifting is explicable in terms of a flexible learning mechanism that learns stronger cue contingencies earlier on, and gradually learns the weaker cue contingencies too. This is supported

by the regression analyses of learning, as well as by the explicit judgments that participants make about the cue-outcome relations and their use of them.

### *Overshooting and Maximizing*

Another benefit of the regression analyses is that they allow us to compare an individual's actual learning profile with that of an ideal learner exposed to the same (possibly) idiosyncratic pattern of information. The main finding was that people's implicit cue weights tracked those of an ideal learner across the course of the task, both with respect to their direction and their strength. However, the analysis also revealed a tendency for participants' implicit weights to overshoot the ideal weights, especially towards the end of the task.

Why do participants' implicit cue weights overshoot the ideal values? One common explanation for this is that people use a maximizing rather than a matching choice rule (Friedman & Massaro, 1998). The main idea here is that once subjects have assessed the evidential weight of the cue pattern on a particular trial, and translated this into an overall weight for each response (fine or rain), they still need to adopt a choice rule. The two main candidates are a maximizing rule, where the response with the greater weight is definitely chosen, or a matching rule, where responses are chosen stochastically in proportion to their relative weights. For example, if a cue pattern gives a probability of 0.75 to outcome A, and 0.25 to outcome B,

maximizing would lead to choosing A every time, whereas matching involves choosing A and B in proportion to these probabilities.

Such a tendency to maximize rather than match would account for the overshooting observed in the current task. This is because the patterns of responses made by a maximizer will look (in terms of the revealed weights) as if they are placing a higher weight on that cue than is justified by its actual weight. This is vividly demonstrated if we construct a learning profile for an ideal maximizer who has pre-knowledge of the objective probabilities for each cue pattern. Their overall regression weights for each cue (computed across all 200 trials) are [-38, -19, +19, +38], which clearly overshoot the objective weights of [-2.1, -0.6, +0.6, +2.1] (which would be achieved by an ideal 'matcher'). So a plausible explanation for the overshooting observed in these experiments is that participants are tending to maximize rather than match, especially towards the end of the task (presumably once they are more sure about the correct responses). Indeed we would predict that if participants were exposed to sufficiently long training their revealed cue weights would tend towards the extreme values of an ideal maximizer.

In what sense is it rational to maximize rather than match in these kinds of task? There has been considerable discussion of the pros and cons of maximizing vs. matching (Friedman & Massaro, 1998; Shanks, Tunney & McCarthy, 2002). Applied to this specific task maximizing is in fact the optimal choice rule. This is because the outcomes are exclusive, and so there is

no informational gain from selecting the less probable option. However, in slightly different environments probability matching might be the better rule. For example, consider an environment in which the outcomes for each response are not exclusive (e.g., finding out about the success of response A tells you nothing about the success of response B). In this case matching may be the best choice rule to adopt until one learns the success rates associated with both responses. And it might remain a good strategy if the environment is non-stationary.

These considerations suggest that there are at least two processing stages in a binary prediction task (Friedman & Massaro, 1998; Kitizis et al., 1998): An evidence stage in which the cue-outcome relations are learned, and a choice stage in which these relations are used to choose a response. Indeed the evidence stage can be further broken down into an acquisition phase, where the separate cue-outcome relations are learned, and an integration phase, where these are combined to yield a weighting for each response option. In the current task the integration stage can be a simple combination rule (e.g., adding the regression weights of the present cues). The choice stage, as we have mentioned, can vary between a maximizing and matching strategy.

This also points to one shortcoming of the regression technique – the cue weight profile by itself cannot distinguish between a maximizer with accurate cue weights and a matcher with over-estimated cue weights.

However, in the current studies the explicit measures of participants' cue weights suggest that roughly halfway through the task they have acquired accurate knowledge of the actual cue weights, and thereafter they tend to use a maximizing choice rule.

#### *Relevance to Neuropsychological Research*

Probabilistic category learning is a fundamental cognitive ability, and one that can be severely impaired through brain dysfunction. Consequently, multicue learning tasks such as the WP task are often used to study cognitive deficits that arise from Parkinson's disease (Knowlton et al., 1996), amnesia (Knowlton, Squire & Gluck, 1994; Hopkins et al., 2004), Huntington's disease (Knowlton et al., 1994) and numerous other neurological conditions. The widespread usage of multicue tasks as a diagnostic tool in patient populations heightens the need to develop a clear understanding of how people master (or fail to master) them. As noted above, our findings suggest that certain standard interpretations of multicue learning are open to question. In particular, we have contested the claims that such tasks engage an implicit procedural learning system (inaccessible to awareness); that people lack insight when learning; that they adopt suboptimal strategies etc. This means that some of the conclusions drawn from previous research with patient groups may require re-evaluation.

This is well illustrated by recent work in the neuropsychological literature that attempts to explain why Parkinson's disease patients perform

poorly on probabilistic learning tasks. The standard explanation for this deficit is that the damage to the Basal Ganglia suffered by Parkinson's patients impairs their procedural learning system (Knowlton et al., 1996). As we have already noted, however, this kind of explanation is questionable. It seems, in normal subjects at least, that multicue learning takes place with full awareness and insight on the part of the subject. In terms of the standard dichotomy this would implicate an explicit learning system. This problem might be circumvented by supposing that Parkinson's patients tackle the task in a very different way from normal subjects. However, a simpler explanation is that these patients suffer from a generalized learning decrement.

Similar issues are raised in a recent study by Shohamy, Myers, Onlaor, and Gluck (2004). They conjecture that Parkinson's patients under-perform on the WP task because they are restricted in the kind of strategies that they can use to solve the task. In their study Parkinson's patients and controls completed the same WP task three times on consecutive days. 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. Indeed their overall performance at the end of day 3 closely resembled the overall performance of controls at the end of day 1.

To account for this pattern of results Shohamy et al. (2004) argued that Parkinson's patients use different strategies to controls in order to learn the task. In particular, they maintained that damage to the Basal Ganglia restricts

patients to singleton or single-cue strategies rather than an optimal multicue strategy. Such strategies enable above chance performance, but place limits on the maximum achievable performance. Their main evidence for this claim was that even by day 3 the majority of patients were best fit by a singleton strategy, whereas the majority of controls, after the same amount of training, were best fit by a multicue strategy.

However, the experimental data are equally well (perhaps better) explained in terms of a generalized learning deficit. After all, learning in Parkinson's patients is known to be slower on a variety of cognitive tasks. And a general decrement in learning would explain their reduced performance relative to controls, as well as the fact that the strategy model fits for Parkinson's patients on day 3 were identical to the model fits for the controls on day 1. The apparent shift from singleton to multicue strategies is again explicable by the gradual learning of each cue-outcome association. As noted above, the early stages of learning may be well fit by a singleton model, but this may reflect imperfect multicue learning rather than a singleton strategy.

On this alternative account, then, the degradation in performance seen in Parkinson's patients results from a general drop in learning rate, not a specific inability to engage in multicue learning. Of course further research is needed to decide definitively between these alternatives, but it is important to acknowledge that a viable, and indeed more parsimonious, alternative exists.

A natural extension would be to test Parkinson's and other patients using the online measures of cue usage introduced in the current study.

### *Conclusions*

In this paper we have shown that multicue probability learning is accompanied by explicit task knowledge and self-insight. Although this accords well with commonsense, it is a position that has been questioned in recent research. In addition we have argued that individual learning profiles are well explained in terms of a general learning mechanism that integrates several cue-outcome contingencies. A dynamic rolling regression model gives a good qualitative fit to the human data, and thus provides us with a convincing as-if model to describe human behaviour. There are a variety of psychologically plausible mechanisms that could instantiate this computational level description (e.g., connectionist networks; Stone, 1986). Finally, we have conjectured that a single declarative learning system is sufficient to explain people's online predictions, their probability judgments, and their explicit reports about their own cue usage. In short, a successful learner is not only sensitive to the statistical structure of the environment, but is sensitive to this sensitivity.

## References

- Ashby, F. G., & Ell, S. W. (2001). The neurobiology of human category learning. *Trends in Cognitive Sciences*, 5, 204–210.
- 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.
- Brunswik, E. (1943). Organismic achievement and environmental probability. *Psychological Review*, 50, 255–72.
- Brehmer, B. (1979). Preliminaries to a psychology of inference. *Scandinavian Journal of Psychology*, 20, 193–210.
- Brehmer, B. (1980). In one word: not from experience. *Acta Psychologica*, 45, 223–241.
- Cooksey, R. (1996). *Judgment analysis: theory, methods, and applications*. Academic Press, London.
- Dayan, P., Kakade, S., & Montague, P.R. (2000). Learning and selective attention. *Nature Neuroscience*, 3, 1218–1223.
- Doherty, M. E., & Kurz, E. (1996). Social judgment theory. *Thinking and Reasoning*, 2, 109–140.
- Doherty, E.D. & Balzer, W.K. (1988). Cognitive feedback. In B. Brehmer & C.R.B. Joyce (Eds.), *Human Judgment: The SJT View*. New York: North-Holland, 115–162.

- Ericsson, K. and Simon, H. (1980). Verbal reports as data. *Psychological Review*, 87, 215–251.
- Evans, J. St. B. T., Clibbens, J., Cattani, A., Harris, A., & Dennis, I. (2003). Explicit and implicit processes in multicue judgment. *Memory & Cognition*, 31, 608-618.
- Friedman, D., Massaro, D. W., Kitzis, S. N., & Cohen, M. M. (1995). A comparison of learning models. *Journal of Mathematical Psychology*, 39, 164-178.
- Friedman, D. & Massaro D.W. (1998). Understanding variability in binary and continuous choice. *Psychonomic Bulletin & Review*, 5, 370–389.
- Gigerenzer, G., Todd, P.M., & the ABC Research Group. (1999). *Simple heuristics that make us smart*. New York: Oxford University Press.
- Gluck, M. A., & Bower, G. H. (1988). From conditioning to category learning: An adaptive network model. *Journal of Experimental Psychology: General*, 117, 227–247.
- 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.
- Goldstein, W. M. (2004). Social Judgment Theory: Applying And Extending Brunswik’s Probabilistic Functionalism. In D. Koehler & N. Harvey (Eds.), *Blackwell handbook of judgment and decision making*. Blackwell.

- Hammond, K. R. (1955). Probabilistic functioning and the clinical method. *Psychological Review*, 62, 255-262.
- Harries, C. & Harvey, N. (2000). Taking advice, using information and knowing what you are doing. *Acta Psychologica*, 104, 399–416.
- Harries, C., Evans, J.St.B.T. and Dennis, I. (2000) Measuring Doctors' Self-Insight into their Treatment Decisions. *Applied Cognitive Psychology*, 14, 455-477.
- Hastie, R., & Dawes, R. M. (2001). *Rational choice in an uncertain world*. Thousand Oaks, CA: Sage.
- 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.
- Kahneman, D., Slovic, P., & Tversky, A. (1982). *Judgment under uncertainty: Heuristics and biases*. Cambridge: Cambridge University Press.
- Kelley, H., & Friedman, D. (2002). Learning to forecast price. *Economic Enquiry*, 40, 556 – 573.
- Kitzis S., Kelley, H., Berg, D., Massaro, D., & Friedman, D. (1998). Broadening the tests of learning models. *Journal of Mathematical Psychology*, 42, 327-55.
- Klayman, J. (1988). On the how and why (not) of learning from outcomes. In B. Brehmer & C. R. B. Joyce (Eds.), *Human Judgment: The SJT view* (pp. 115-160). North-Holland: Elsevier.

- Knowlton, B., Mangels, J., & Squire, L. (1996). A neostriatal habit learning system in humans. *Science*, *273*, 1399 – 1402.
- Knowlton, B., Squire, L. and Gluck, M. (1994). Probabilistic classification learning in amnesia. *Learning & Memory*, *1*, 106 – 120.
- 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.
- Myers, C. E., Shohamy, D., Gluck, M. A., Grossman, S., Kluger, A., & Ferris, S. (2003). Dissociating hippocampal vs. basal ganglia contributions to memory using a two-phase test of learning and transfer. *Journal of Cognitive Neuroscience*, *15*(2), 185–193.
- Nisbett, R. and Wilson, T., 1977. Telling more than we can know: Verbal reports on mental processes. *Psychological Review* *84*, 231–259.
- Poldrack, R. A., Clark, J., Pare-Blagojev, E. J., Shohamy, D., Creso Moyano, J., Myers, C., & Gluck, M. A. (2001). Interactive memory systems in the human brain. *Nature*, *414*, 546–550.
- 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.
- Reilly, B., & Doherty, M. (1992). The assessment of self-insight judgment policies. *Organizational Behavior and Human Performance*, *53*, 285–309.

- Shanks, D. R. (1990). Connectionism and the learning of probabilistic concepts. *Quarterly Journal of Experimental Psychology*, 42A, 209–237.
- Shanks, D. R., & St John, M. F. (1994). Characteristics of dissociable human learning systems. *Behavioural and Brain Sciences*, 17, 367–447.
- 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. (2004). Role of the Basal Ganglia in category learning: How do patients with Parkinson's disease learn? *Behavioral Neuroscience*, 118, 4, 676–686.
- Slovic, P. and Lichtenstein, S. (1971). Comparison of Bayesian and regression approaches to the study of information processing in judgment. *Organizational Behavior and Human Performance*, 6, 649–744.
- Smith, A. C., Frank, L. M., Wirth, S., Yanike, M., Hu, D., Kubota, Y., Graybiel, A. M., Suzuki, W. A., Brown, E. N. (2004). Dynamic Analysis of Learning in Behavioral Experiments. *Journal of Neuroscience*, 24, 447-461.
- Stone, G. O. (1986). An analysis of the delta rule and the learning of statistical associations. In D. E. Rumelhart, J.L. McClelland & the PDP Research Group (Eds.), *Parallel distributed processing: explorations in the microstructure of cognition* (pp. 444–459), 1. Cambridge, MA: MIT Press.

- Squire, L. R. (1994). Declarative and nondeclarative memory: Multiple brain systems supporting learning and memory. In D. L. Schacter & E. Tulving (Eds.), *Memory systems* (pp. 203–231). Cambridge, MA: MIT Press.
- White, P. A. (1988). Knowing more than we can tell: 'Introspective access' and causal report accuracy 10 years later. *British Journal of Psychology*, 79(1), 13–45.
- Wigton, R.S. (1996). Social judgment theory and medical judgment. *Thinking and Reasoning*, 2, 175-190.
- York, K. M, Doherty, M. E., & Kamouri, J. (1987). The influence of cue unreliability on judgment in a multiple cue probability learning task. *Organizational Behavior and Human Performance*, 39, 303 – 317.

Table 1

*Learning Environment for Experiments 1 & 2.*

Pattern	Cards present	Total	P(pattern)	P(fine pattern)
A	0001	19	0.095	0.895
B	0010	9	0.045	0.778
C	0011	26	0.13	0.923
D	0100	9	0.045	0.222
E	0101	12	0.06	0.833
F	0110	6	0.03	0.500
G	0111	19	0.095	0.895
H	1000	19	0.095	0.105
I	1001	6	0.03	0.500
J	1010	12	0.06	0.167
K	1011	9	0.045	0.556
L	1100	26	0.13	0.077
M	1101	9	0.045	0.444
N	1110	19	0.095	0.105
Total		200	1.00	

Note: 0 = card present, 1 = card absent.

## Figure Captions

*Figure 1.* Probabilistic environment in the Weather Prediction task for Experiments 1 & 2.

*Figure 2.* The Lens Model: the objective weights from cues and outcome ( $w_1, w_2, w_3, w_4$ ) are supposed to parallel the subjective weights from cues to judgment ( $j_1, j_2, j_3, j_4$ ).

*Figure 3.* Screen presented to participants in the learning phase of Experiments 1 & 2.

*Figure 4.* Learning performance measured by mean proportion correct predictions ( $\pm$ SE) in Experiment 1 (top panel) and Experiment 2 (bottom panel). In both experiments the optimal level is 0.83.

*Figure 5.* Blocked mean probability ratings ( $\pm$ SE) for rain given each card in Experiment 1 (top panel) and Experiment 2 (bottom panel). The objective probabilities for each card are given by the dotted lines (card 1 = 20%, card 2 = 40%, card 3 = 60%, card 4 = 80%).

*Figure 6.* Blocked importance ratings ( $\pm$ SE) for strong and weak cards in Experiment 1 (top panel) and Experiment 2 (bottom panel).

*Figure 7.* Individual judgment profiles for the strong cards (1 and 4) for two participants in Experiment 1 (top and bottom panels). Card 4 is strongly predictive of sun (+ regression weight) and card 1 is strongly predictive of rain (- regression weight). The solid lines are the participant's subjective judgments weights, the dashed lines are the weights of an ideal judge exposed to the same environment.

*Figure 8.* Averaged judgment profiles for the two strong cards for all participants in Experiment 1 (top panel) and Experiment 2 (bottom panel).

*Figure 9.* Averaged judgment profiles for the two weak cards for all participants in Experiment 1 (top panel) and Experiment 2 (bottom panel).

*Figure 10.* Averaged judgment weights (absolute values) for the strong cards (1&4) vs. weak cards (2&3) in Experiment 1 (top panel) and Experiment 2 (bottom panel).

*Figure 11.* Mean trial-by-trial importance ratings in Experiment 2. The top panel shows ratings across individual trials, the bottom panel ratings averaged across blocks of ten trials. Note that values on the y-axis represent how much participants said they relied on each card, with 4 = "Greatly", 3 = "Moderately", 2 = "Slightly", 1 = "Not at all".

*Figure 12.* Individual judgment profiles for the strong cards (1 and 4) for two participants in Experiment 2 (top and bottom panels). Card 4 is strongly predictive of sun (+ regression weight) and card 1 is strongly predictive of rain (- regression weight).

*Figure 13.* Two possible models of the relation between learning and behavioural responses. A: single process model; B: dual-process model (adapted from Lovibond and Shanks, 2002).

Figure 1.

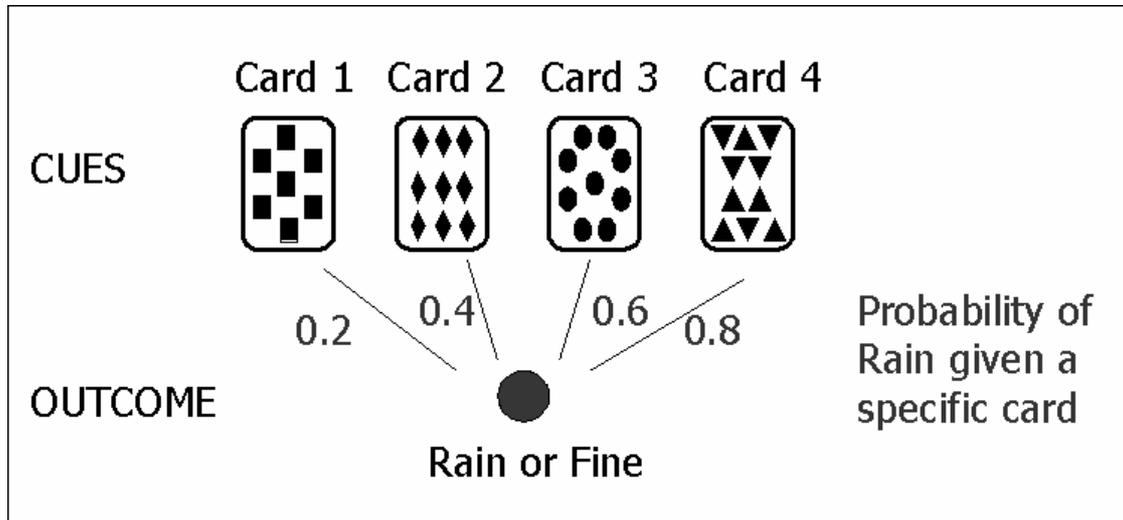


Figure 2.

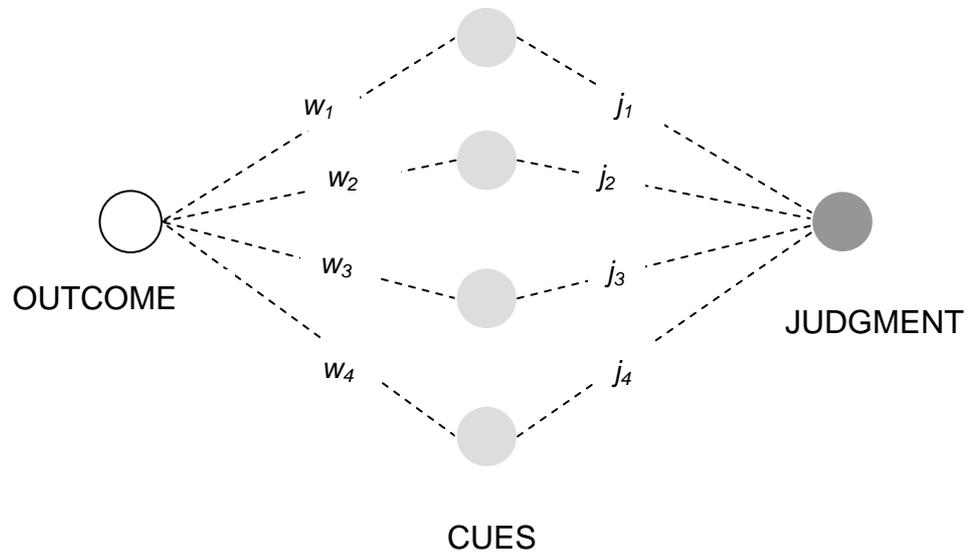
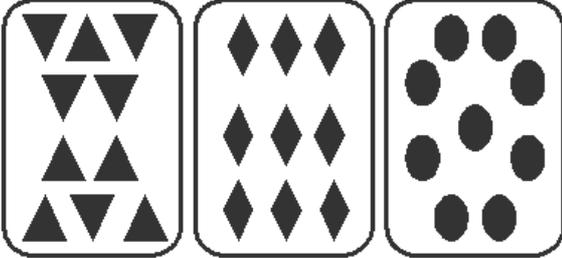


Figure 3.

What is the weather going to be like? Rain or Fine?



**RAINY**

**FINE**



Earnings

Figure 4.

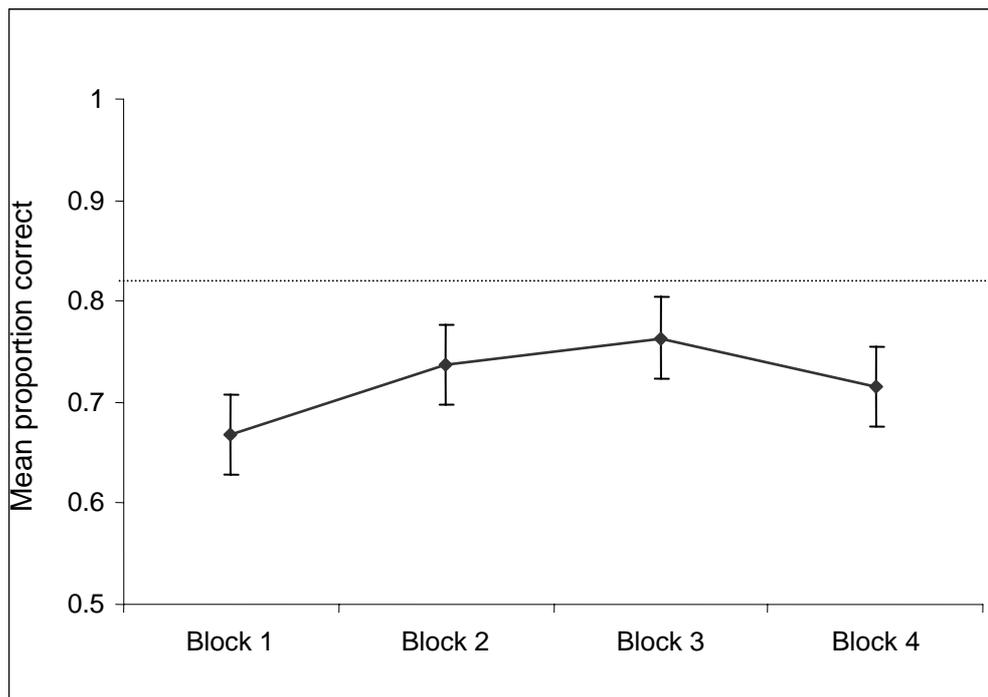
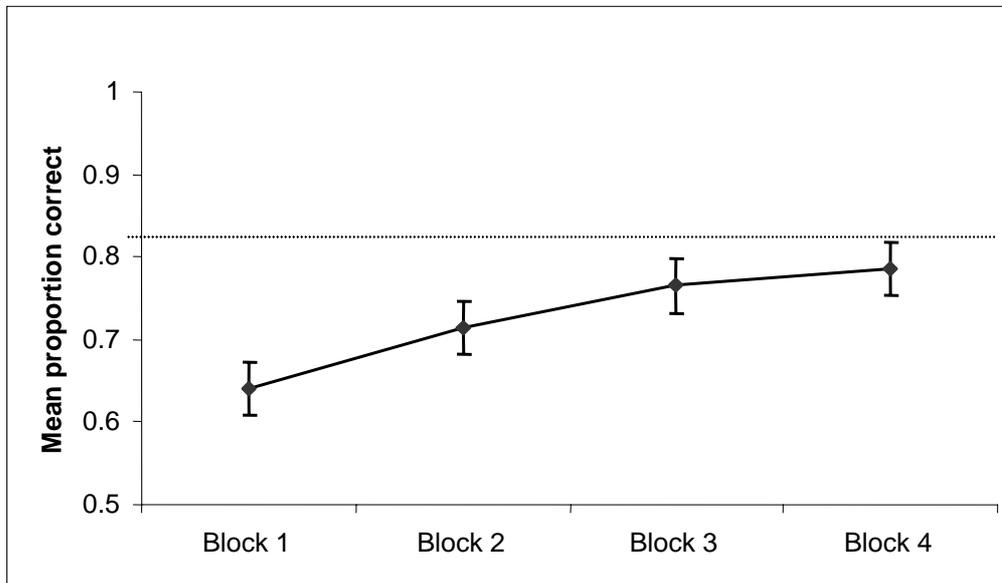


Figure 5.

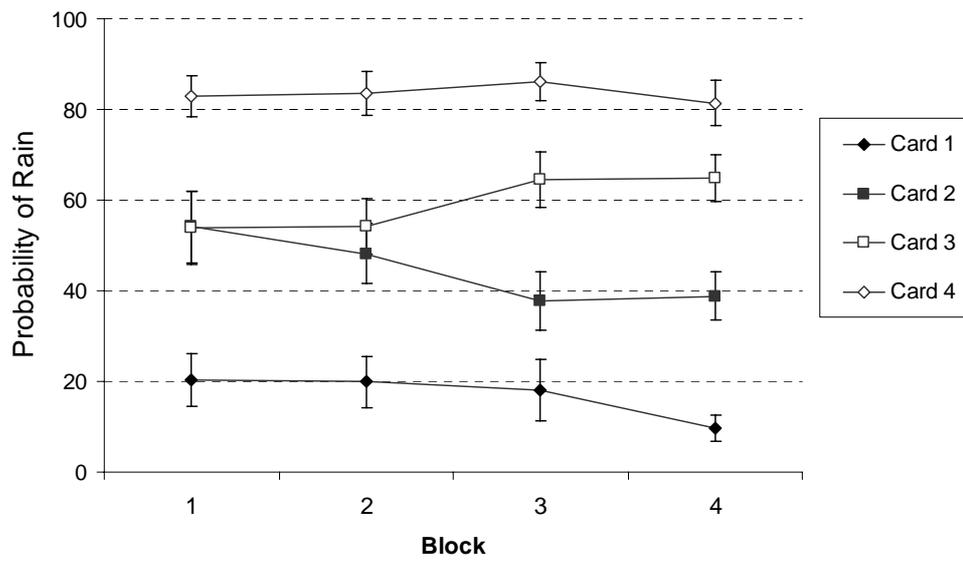
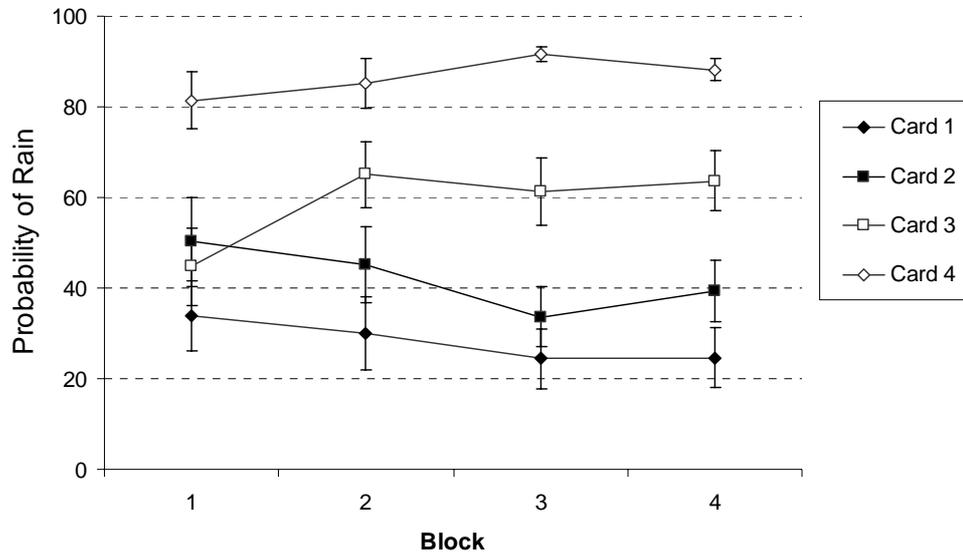


Figure 6.

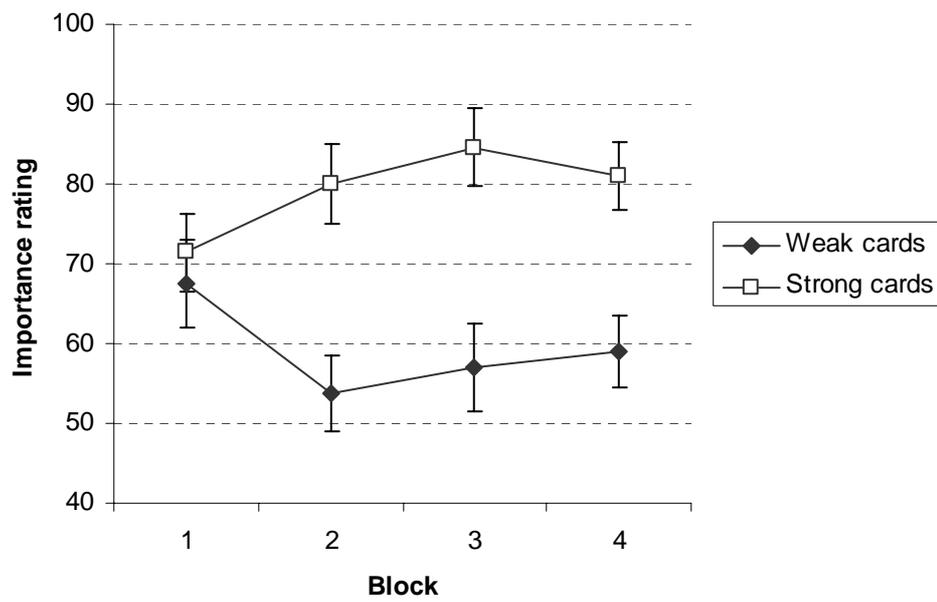
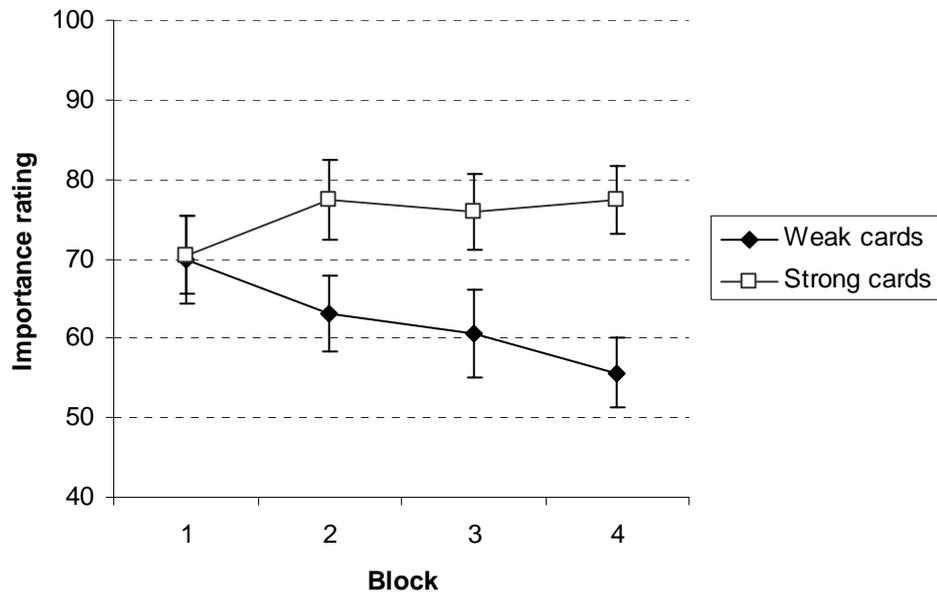


Figure 7.

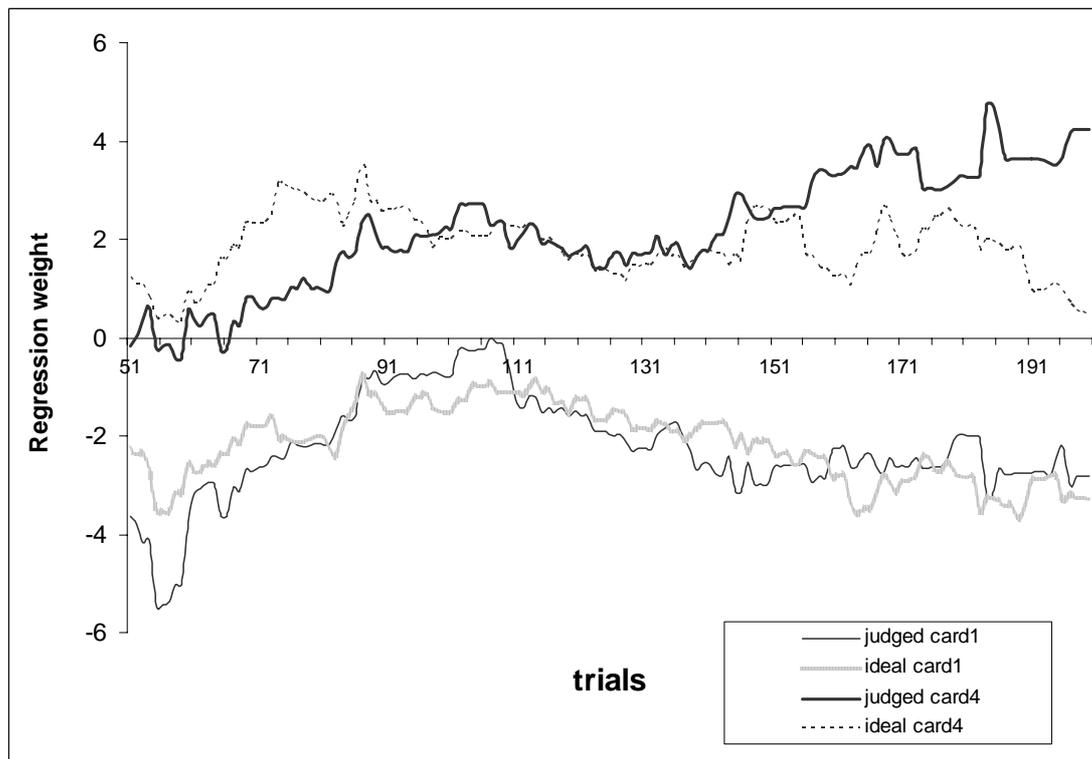
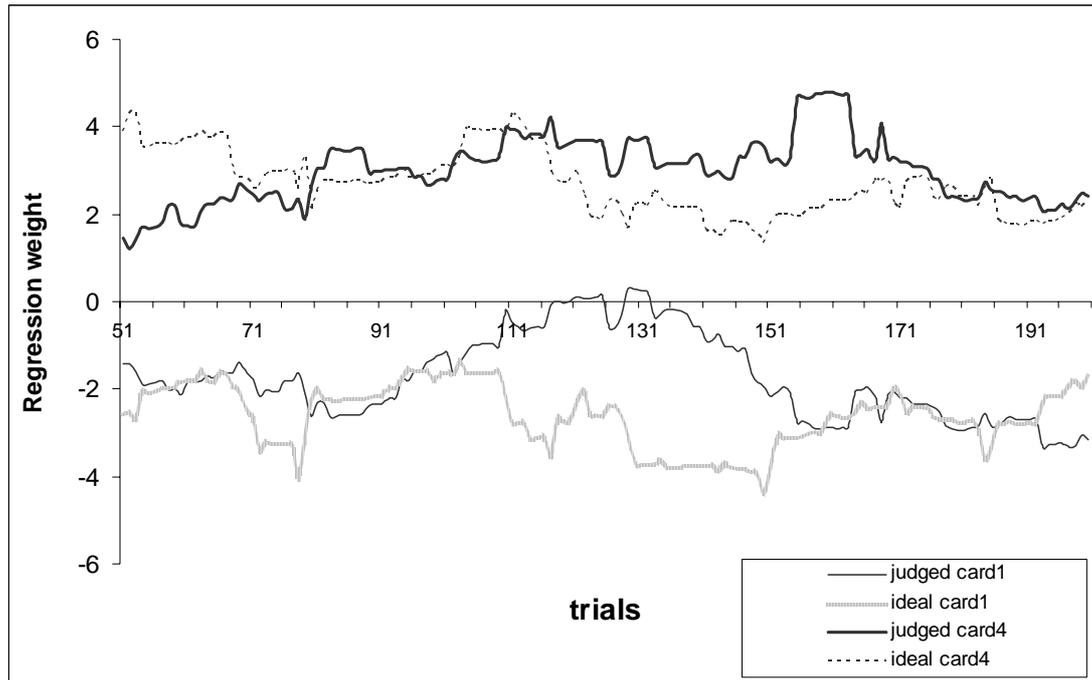


Figure 8.

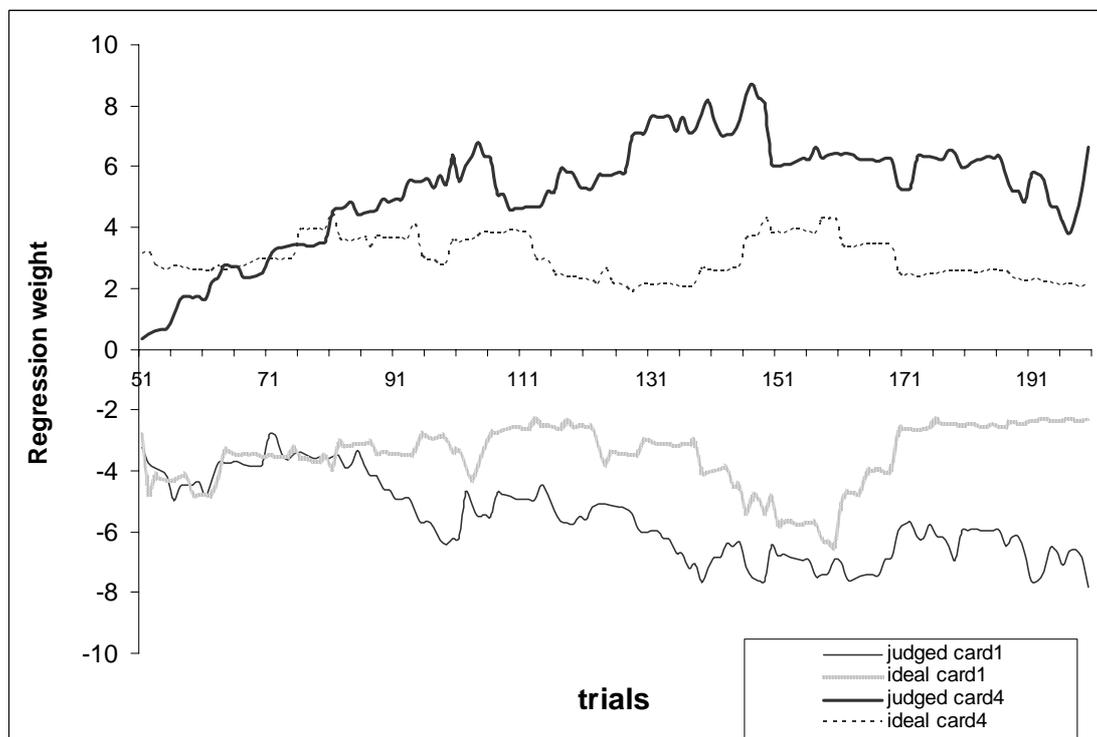
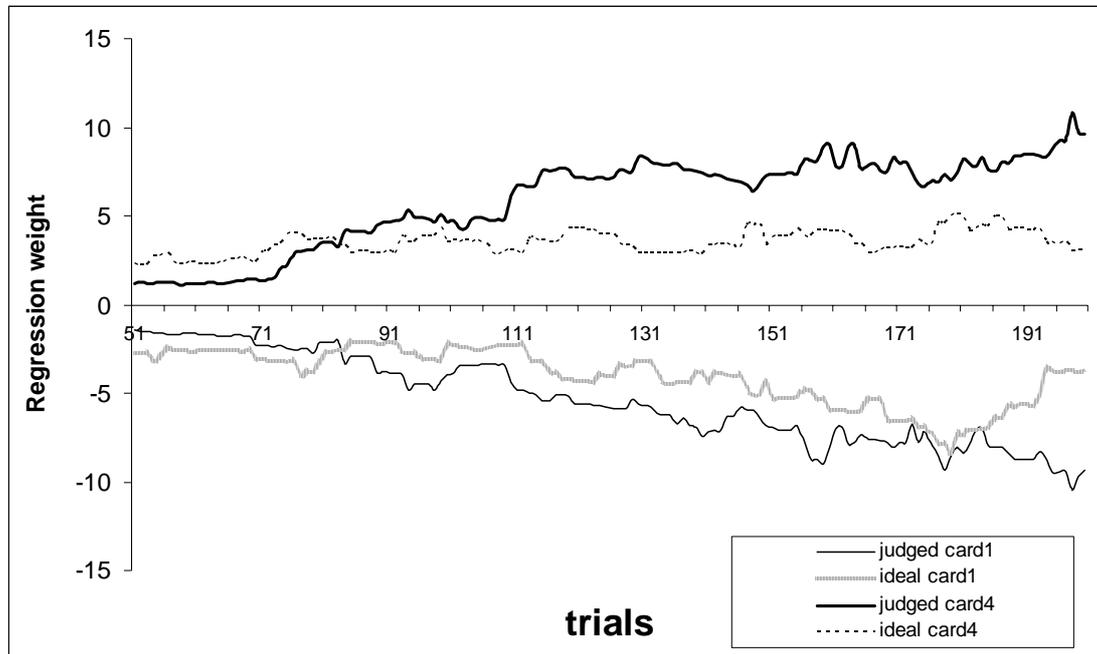


Figure 9.

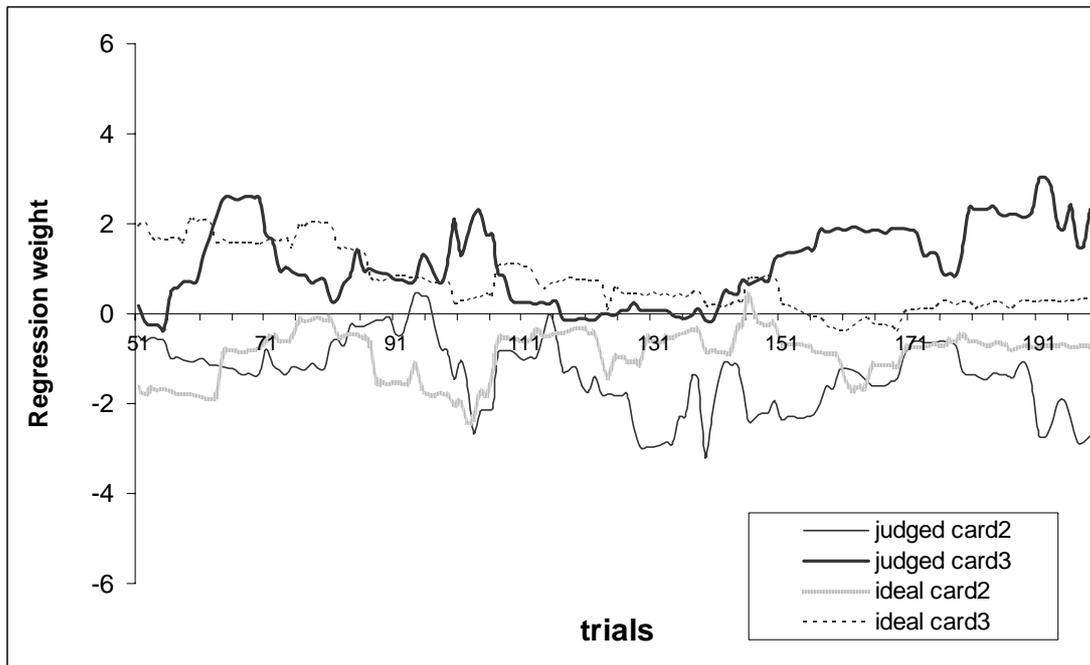
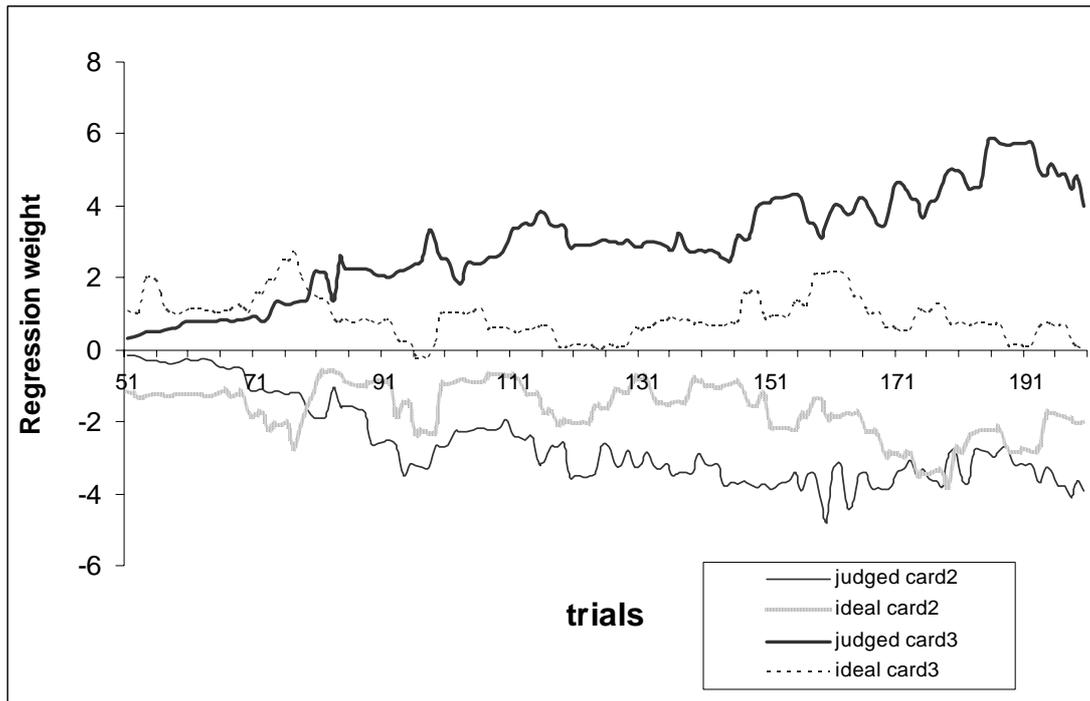


Figure 10.

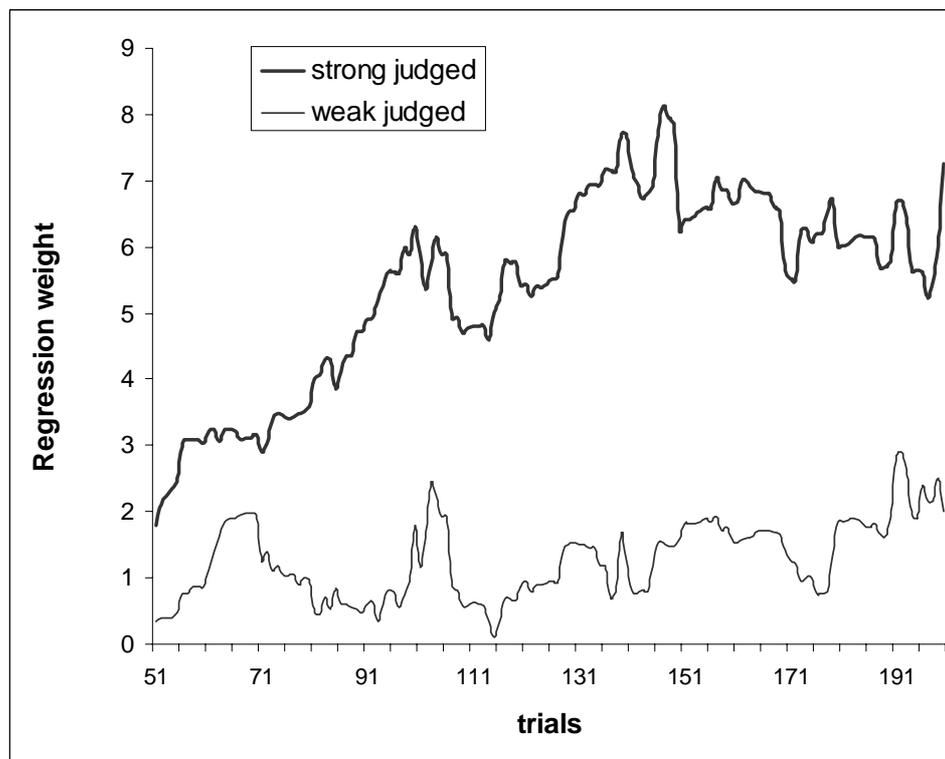
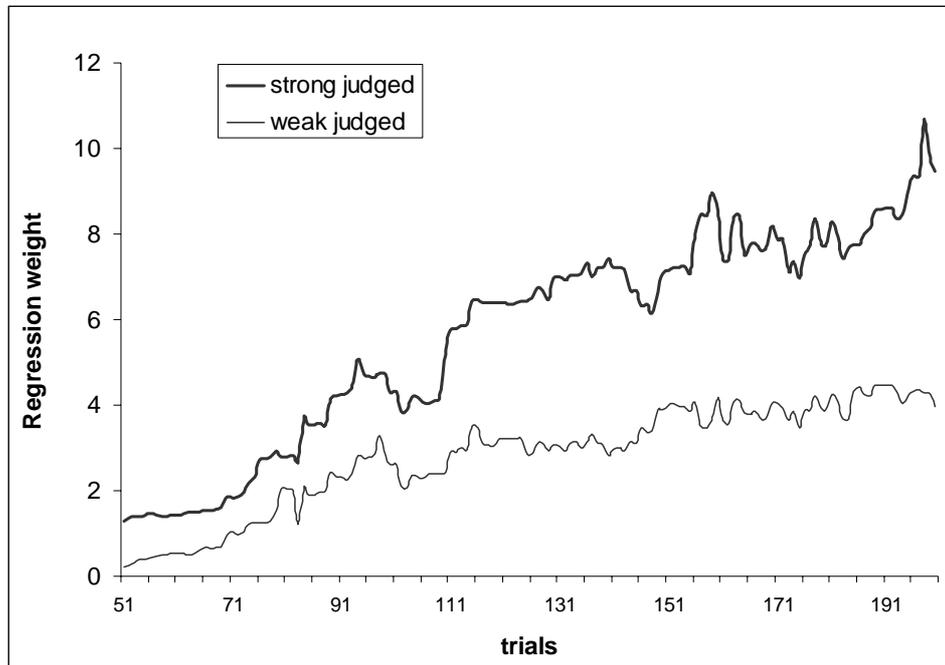


Figure 11.

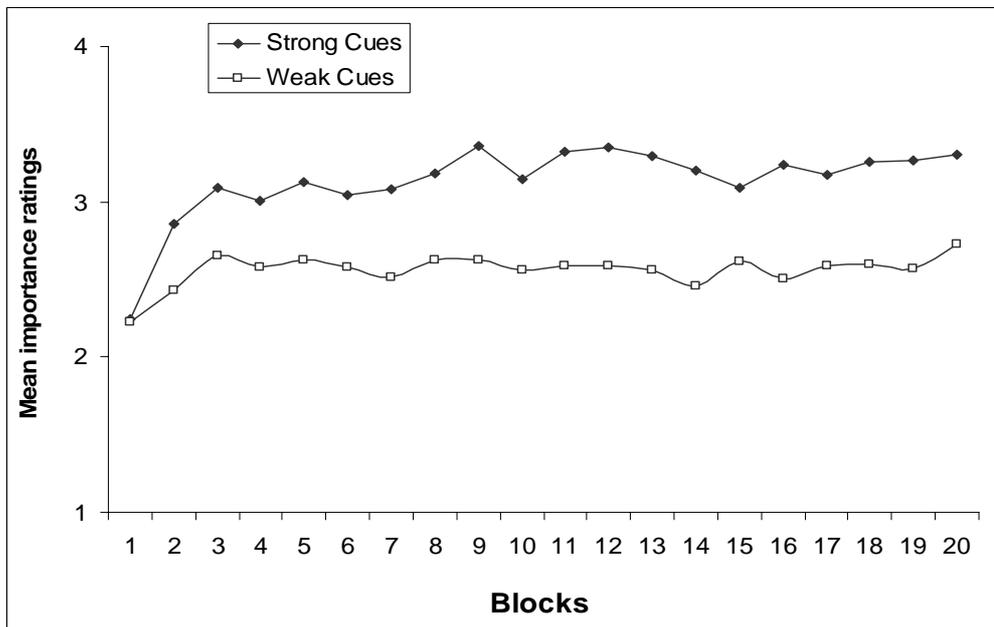
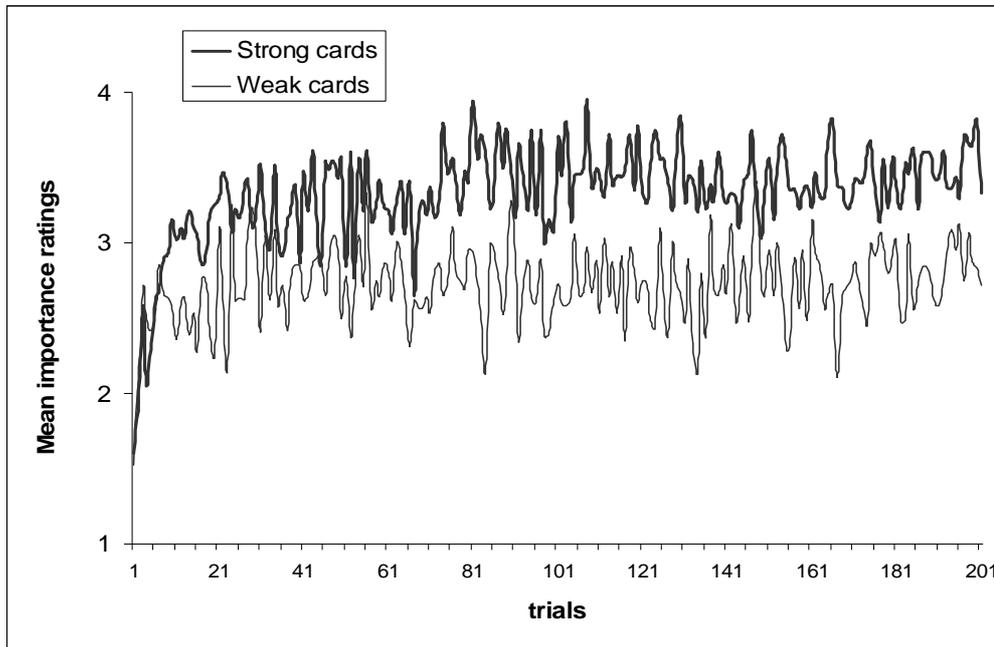


Figure 12.

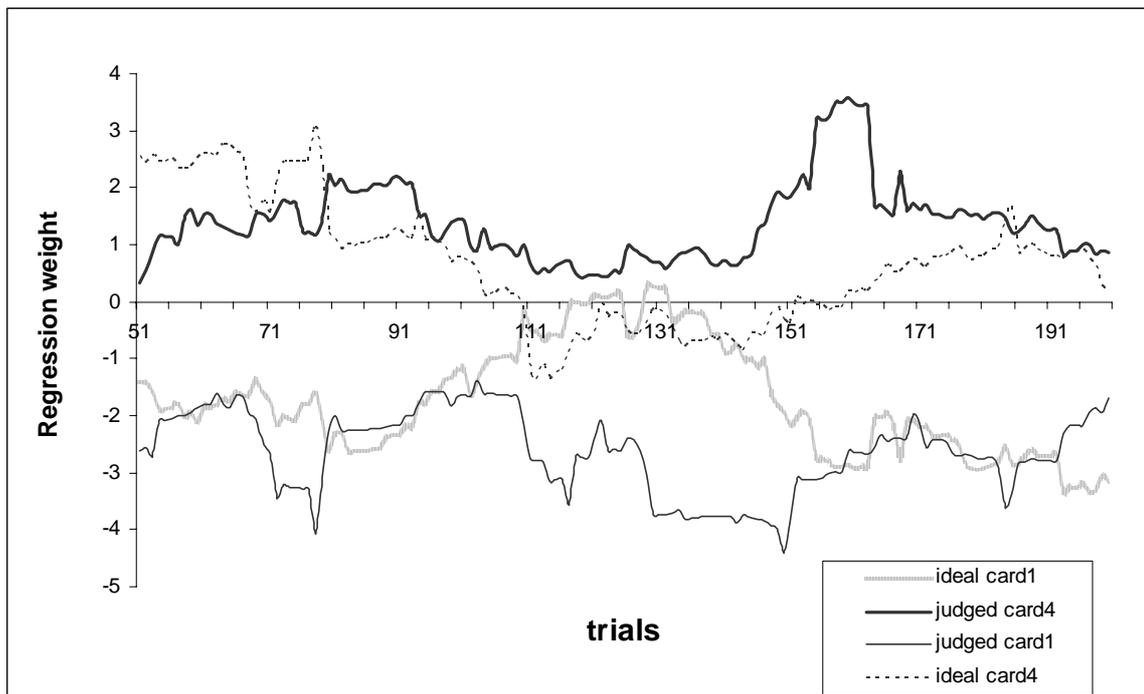
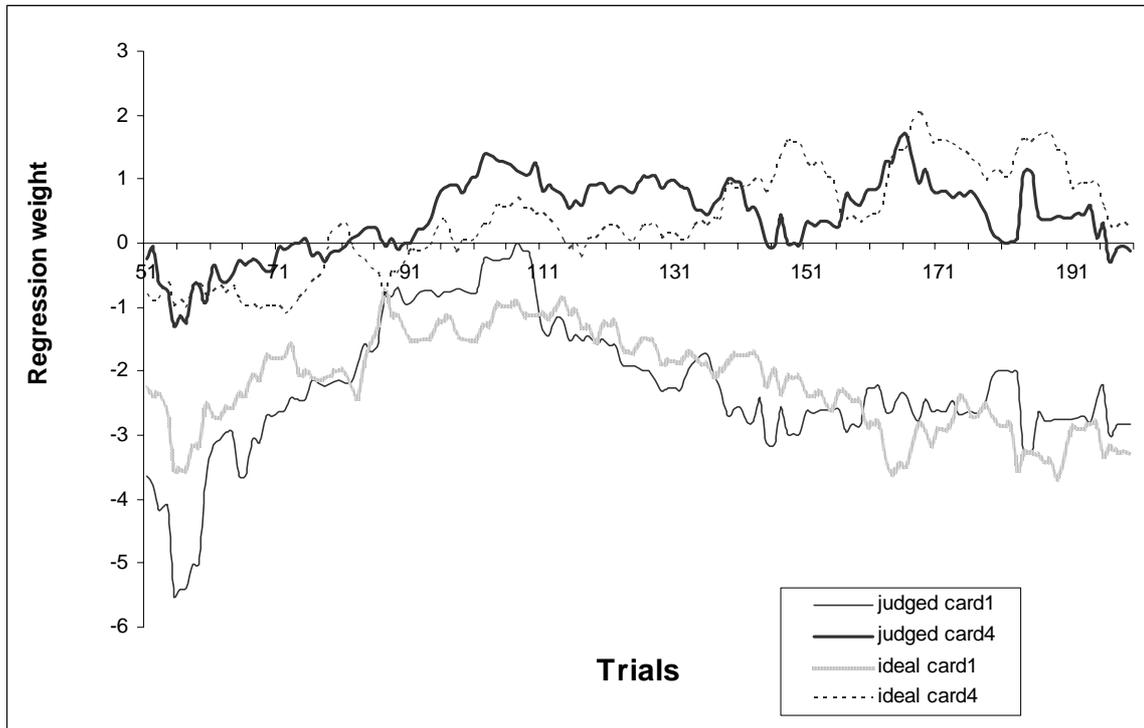
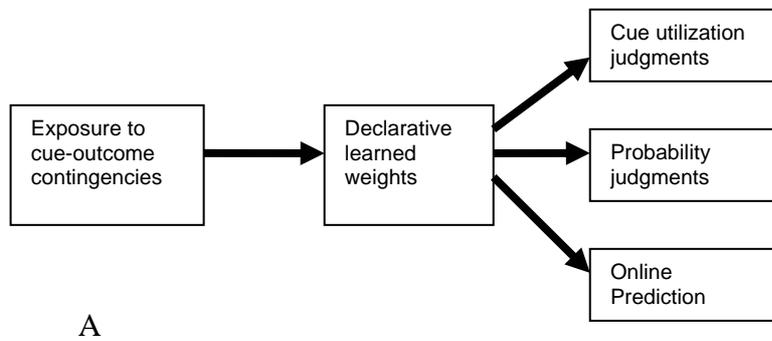
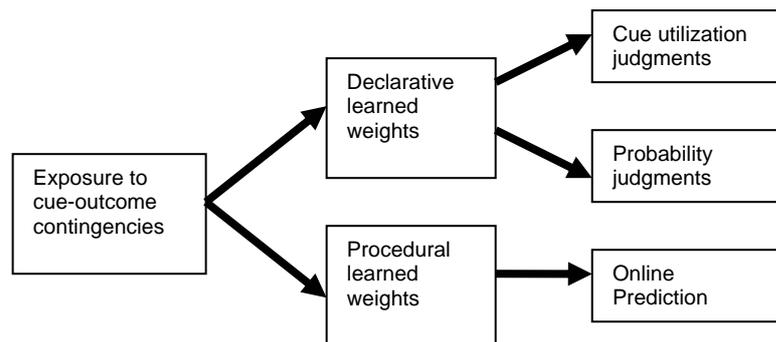


Figure 13.



A



B