

Learning in a changing environment

Maarten Speekenbrink and David R. Shanks

Division of Psychology and Language Sciences
University College London

Draft of December 10, 2009

Abstract

Multiple cue probability learning studies have typically focused on stationary environments. We present three experiments investigating learning in changing environments. A fine-grained analysis of the learning dynamics shows that participants were responsive to both abrupt and gradual changes in cue-outcome relations. We found no evidence that participants adapted to these types of change in qualitatively different ways. Also, in contrast to earlier claims that these tasks are learned implicitly, participants showed good insight into what they learned. By fitting formal learning models, we investigated whether participants learned global functional relationships or made localized predictions from similar experienced exemplars. Both a local (the Associative Learning Model) and a global learning model (the novel Bayesian Linear Filter) fitted the data of the first two experiments. However, the results of Experiment 3, which was specifically designed to discriminate between local and global learning models, provided more support for global learning models. Finally, we present a novel model to account for the cue competition effects found in previous research and displayed by some of our participants.

Nothing endures but change

Heraclitus (540 BC – 480 BC).

Predicting future events from past experience is a fundamental aspect of daily life. For instance, policy makers have to predict the outcome of interventions and sports coaches must

The support from the Economic and Social Research Council is gratefully acknowledged. This research was part of the ESRC Centre for Economic Learning and Social Evolution (ELSE) and partially supported by ESRC grant RES-062-23-1511.

The authors would like to thank Jerome Busemeyer, Peter Dayan and three anonymous reviewers for helpful comments on previous versions of the manuscript.

Address correspondence to M. Speekenbrink, Cognitive Perceptual and Brain Sciences, Division of Psychology and Language Sciences, University College London, 26 Bedford Way, London WC1H 0AP, England, e-mail: m.speekenbrink@ucl.ac.uk

predict a player's influence on team performance. But as many environments are subject to complex changes over time, these tasks can be challenging. In times of economic instability, previously successful interventions may not have the desired effect. And a player's past performance may no longer be predictive when (s)he loses form. Not only are the cues informing predictions noisy, their predictive validity may change as the environment evolves.

This article addresses how people learn to predict outcomes from multiple cues when the relations between them undergo unsignalled changes. While a relevant and important topic, surprisingly little research has addressed it. Previous research has mainly considered stable environments. And compared to category learning, little attention has been paid to how people learn to predict metric outcomes (Busemeyer, Byun, DeLosh, & McDaniel, 1997; DeLosh, Busemeyer, & McDaniel, 1997). Although this situation is slowly changing (Busemeyer et al., 1997; DeLosh et al., 1997; Kalish, Lewandowsky, & Kruschke, 2004), this work usually concerns learning simple functions of a single variable, not learning from multiple cues (but see Juslin, Olsson, & Olsson, 2003; Kelley & Friedman, 2002; Kelley & Busemeyer, 2008). The latter Multiple Cue Probability Learning (MCPL) tasks are popular in judgement research (e.g., Brehmer, 1994; Cooksey, 1996), but this area is usually not directly concerned with learning. The present research attempts to fill this lacuna and in doing so will focus on three main issues: Is learning affected by the abruptness of changes? Do people have insight into these changes? And, is behaviour in multiple cue tasks subserved by a global or a local learning process? While the latter two issues have been addressed previously, it will become clear that the present context of changing environments sheds new light on these issues.

Learning in a changing environment

The ability to adapt to changes in our environment is a vital part of our functioning. But, as the signals in the environment are inherently noisy, how do we distinguish real change from random fluctuation? The problem is one of balance. If we infer too much change, we can't benefit from previous knowledge. Infer too little, and we are stuck in the past. The right balance depends on the relative magnitude of two sources of variation: observational noise (variation of observations from their true values) and structural change (variation in the true values themselves). When the former is relatively large, little change needs to be inferred, as apparent change is most likely observational noise. More change should be inferred if structural change is relatively large, as deviations from expected values are more likely to be due to true change. Research on basic processes such as motor control has shown that humans can perform this balancing act rather well (Baddeley, Ingram, & Miall, 2003; Körding & Wolpert, 2004; Wolpert, Ghahramani, & Jordan, 1995). A much-studied example concerns prism adaptation of visually guided reaching, in which shifted visual feedback drives rapid recalibration of motor movements (Baddeley et al., 2003; Cheng & Sabes, 2007; Chhabra & Jacobs, 2006; Körding & Wolpert, 2004). Higher cognitive abilities such as concept learning have also been shown to be receptive to changes in the environmental structure. For example, normal people perform quite well in the Wisconsin Card Sorting Task (Grant & Berg, 1948; Milner, 1963), which consists of inferring abstract rules that change abruptly once the rule has been learned.

Whether one expects the environment to change or remain stable should affect inference. Consider for example the problem of determining the distance of an object. Numerous cues are used in depth perception, such as interposition, relative height, texture gradient and binocular disparity (e.g., Mather, 2006). For the moment we shall take for granted that the observer

somehow integrates these into a single “depth datum” (an overall immediate perception of the object’s depth). As for each individual cue, the depth datum is a fallible indicator of the object’s true depth. If the object remains stationary, the observer can increase the accuracy of her inference by making multiple observations and averaging these to form an overall estimate. But this is not the best strategy for a moving object, for which recent data are more informative than earlier ones. Rather than a simple average, an optimal observer should form a weighted average, weighting observations according to their recency. The optimal (Bayesian) method for tracking a moving object from noisy data is known as the Kalman filter (Kalman, 1960; Kalman & Bucy, 1961). The Kalman filter implements a method of online Bayesian inference, sequentially updating estimates after each new observation. The Kalman filter is a general statistical tool and the object of estimation can take many forms, such as an object’s depth, but also a weight describing the effect of a cue on a criterion variable. In fact, we use the Kalman filter for precisely this purpose in our Dynamic Lens Model analysis described later. In the remainder of this section, we provide an intuitive description of Bayesian inference – as implemented by the Kalman filter – for our simplified example of depth perception.

We noted that in order to accurately infer change, two sources of uncertainty must be taken into account: the noise in the data (the discrepancy between observed and true values) and structural change (variations in the true values over time). But the observer will usually have some idea about the object’s previous distance, which introduces a third source of uncertainty, related to the accuracy of this prior knowledge. Together, these three sources of uncertainty determine the best response to a new datum. The steps in the Bayesian inference are depicted in Figure 1, both for a moving and stationary object. Bayesian inference concerns updating probability distributions. All prior knowledge of the object’s depth is contained in the the prior distribution (Figure 1A). It represents, for each possible depth value, the observer’s subjective belief that the true depth of the object is identical to that value. When the observer has strong prior knowledge, the prior distribution will have a small variance and be concentrated around a particular value. If the observer’s prior knowledge is weak, the prior distribution will have a large variance and be spread out over a large range of possible values. The current depth of a moving object is likely to differ from its prior depth. Knowledge of the speed and direction in which the object moves can be used to form a predictive distribution of the object’s current depth (Figure 1B). In general, this predictive distribution will have a larger level of uncertainty than the prior distribution. Only when the object is stationary, and there is no uncertainty regarding the object’s movement, will the predictive distribution be identical to the prior distribution. When the observer acquires a new datum (Figure 1C), the observer can use this to update the predictive distribution. Inference would be trivial if the datum was noiseless, as the object’s true depth would simply be identical to the datum. But the datum is unreliable, which is reflected by the variance of the observation distribution – the larger the observational noise, the larger the variance. For noisy data, the observer should take his or her expectation, reflected by the predictive distribution, into account. The relative weight given to the expectation and datum depends on the ratio of the variance of the predictive and observation distribution. The resulting final estimate of the object’s current depth is given by the posterior distribution (Figure 1D). Here, the posterior mean is a weighted average of the means of the predictive and observation distribution. This posterior distribution will take the role of prior distribution at the next time point.

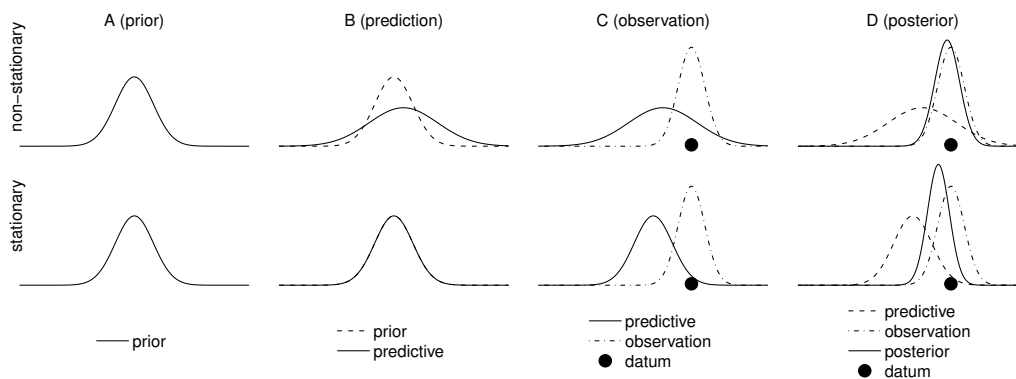


Figure 1. Bayesian inference of a non-stationary (top panels) and stationary (bottom panels) target. The inferential steps are depicted from left to right. In each panel, solid curves represent the current distribution, while dashed curves represent the distribution at the previous stage. Inference starts with a prior distribution (A), which usually reflects the inference of the target at the previous time point. This prior distribution forms the basis for the predictive distribution (B). For a non-stationary target, the predictive distribution differs from the prior distribution due to the uncertainty associated with the target's movement. For a stationary target, there is no additional uncertainty and the predictive distribution is identical to the prior distribution. An observation (solid dot in C) is an imperfect indicator of the target. This is reflected in the observational distribution (dash-dotted line in C). The observation and the predictive distribution form the basis of the posterior distribution (D), which combines the predictive and observation distribution. The posterior distribution functions as the prior distribution at the next time point.

Multiple cue learning

This account of Bayesian inference is simplified in that we took for granted that the observer somehow arrived at a single datum of the object's distance. In reality, this datum itself is formed on the basis of multiple cues, which brings the added complication of how to integrate these cues. Cue integration is a ubiquitous aspect of daily life. For instance, when a stock broker has to predict share price, (s)he should not only take the past share price into account, but also market indicators such as the Dow Jones index, exchange rates, etc. To make accurate predictions, the broker must utilize these cues in accordance with their predictive relation with share price. In other words, the broker must adapt cue utilization to cue validity.

Multiple cue learning has a long history in psychology. Research on this topic usually assumes that participants form predictions of a criterion by an additive combination of cue values (Anderson, 1981; Brehmer, 1994; Cooksey, 1996; Juslin, Karlsson, & Olsson, 2008). That is, participants multiply each cue value by a scaling factor and then add the resulting values to form a final prediction of the criterion value. As such, predictions are a linear function of the cue values, and the effect of each cue on the resulting prediction is independent of the effect of the other cues. The additivity assumption is reflected in the Lens Model (e.g., Cooksey, 1996; Hammond, Wilkins, & Todd, 1966), which was originally proposed by Brunswik (1955) to describe how organisms perceive a distal (unobservable) criterion through a "lens" of proximal (observable) cues. Hammond et al. (1966) formulated a statistical version of the model as a general framework for judgement analysis. The lens model is depicted in the box on the right hand side

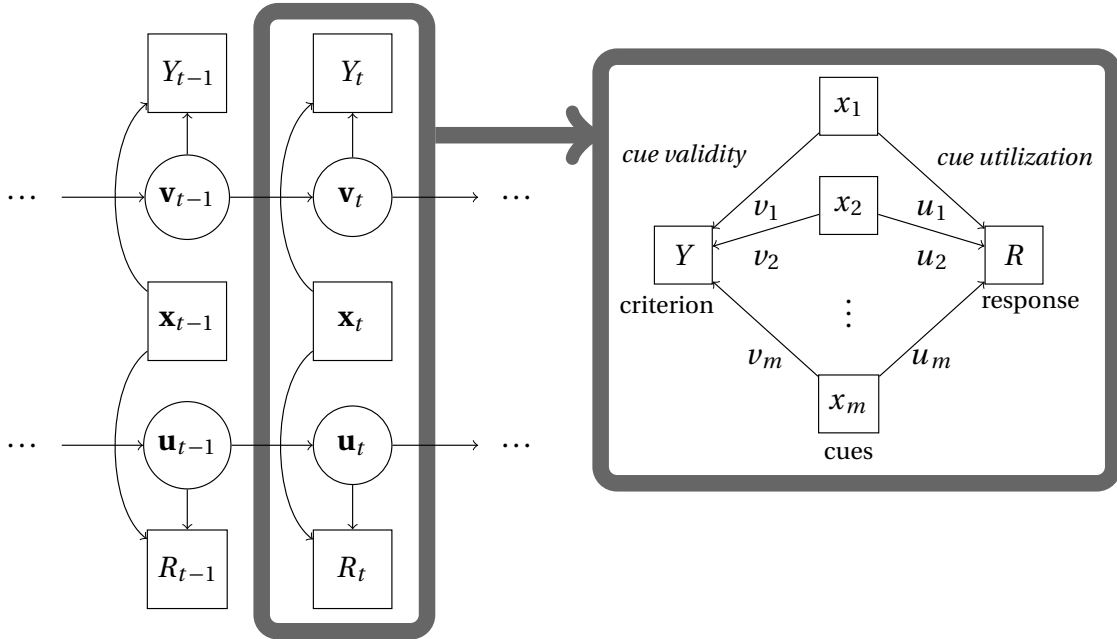


Figure 2. The Dynamic Lens Model (DLM) and classic Lens Model. The DLM (left hand side) consists of two dynamic linear models. In the model of the environment, the criterion values, Y_t , depend on the current (vector of) cue values, \mathbf{x}_t and (vector of) cue validity weights, $\mathbf{v}_t = (v_{1t}, v_{2t})$. In the model of the judge, responses, R_t , depend on the current cue values, \mathbf{x}_t , and (vector of) cue utilization weights, \mathbf{u}_t . Both the cue validity and utilization weights follow a first-order Markov process, such that the validity/utilization at time t depends only on the validity/utilization at time $t - 1$. At a single time point t , the model is equivalent to the classic (static) Lens Model, which is depicted in the box on the right hand side. For comparison to other work, we have depicted the classic Lens model in its standard form, in which cue validity weights, v_j , and utilization weights, u_j , appear as labels of the cue-criterion and cue-response relations, rather than as separate variables as in the dynamic lens model; this is only a superficial difference. For simplicity, there are no time subscripts in the depiction of the classic Lens Model.

of Figure 2. It consists of two linear regression models: one for the effect of the cues on the criterion and one for the effect of the cues on the predictions. The regression coefficients of the first model reflect the structure of the environment and are referred to as cue validity weights. Cue validity weights are optimal scaling factors, those that result in the best linear prediction of the criterion values. The regression coefficients of the second model reflect a participant's prediction policy and are referred to as cue utilization weights. The cue utilization weights are the subjective scaling factors as applied by participants and need not be optimal. By comparing cue utilization to cue validity weights, we can assess how well a participant's prediction policy is adapted to the structure of the environment.

In standard form, the lens model is static; validity and utilization weights are not assumed to change over time. Multiple cue learning research has almost exclusively been conducted with such stationary environments. Exceptions are studies by Dudycha, Dumoff, and Dudycha (1973), Lindberg and Brehmer (1976), Peterson, Hammond, and Summers (1965), Ruffner and Muchinsky (1978), and Summers (1969). These studies have shown that participants are generally re-

sponsive to changes in cue validity but do not optimally adapt their utilization to these changes. For example, in the study by Peterson et al. (1965), the validities of two out of three cues interchanged halfway during the task. The average cue utilization coefficients changed accordingly, but rather late (60 trials after the change in cue validity). Summers (1969) used a three-cue task in which only one cue was predictive (i.e., had a validity weight different from 0). The other two cues were unrelated to the criterion. After 100 trials the validity of the predictive cue either did not change, changed sign (rule shift; if the cue was positively related to the criterion, it would become negatively related, and vice versa), was transferred to another cue (cue shift; one of the previously unpredictable cues took over the role of the predictive cue), or was transferred and changed sign (complete shift). Results showed that participants' performance was best when there was no change, better under rule than cue shift, and worst under a complete shift. Taken together, these previous studies showed that individuals did learn to adapt their utilization to changes in validity, although the speed of this adaptation was not as quick as it should have been. Moreover, some types of change may be more difficult to learn than others. This latter effect has been more extensively investigated in subsequent studies, showing that adaptation to changes in cue validity depends on the particulars of the environment before the change, such as the previous validity of the changing cue (Dudycha et al., 1973; Lindberg & Brehmer, 1976) and the distribution of cue validities over the cues (Ruffner & Muchinsky, 1978).

Dynamic lens model analysis

One problem with the aforementioned studies is that the analyses employed were not particularly responsive to dynamic changes in cue utilization. The analyses consisted of fitting the lens model separately to consecutive blocks of trials. This method is valid insofar as the relation between cues and responses is stable within each block. However, cue utilization will most likely change from trial to trial, such that the estimates of cue utilization may reflect the average utilization in a block, but not the utilization at any single trial. An alternative analysis, rolling regression (e.g., Kelley & Friedman, 2002; Lagnado, Newell, Kahan, & Shanks, 2006), fits regression models to a moving window of trials. This has the advantage of estimating cue utilization weights for each trial after an initialization period. But, although providing a more fine-grained overview of changes in utilization, the method is essentially similar to the simple blocked version and suffers from the same drawbacks. To alleviate these problems, the change process should be explicitly incorporated in the model. We have proposed to do this in a state-space model formalism (Speekenbrink & Shanks, 2008). The resulting dynamic lens model (DLM) consists of two dynamic linear models (e.g. West & Harrison, 1997). Relying on advanced estimation procedures, the DLM provides a fine-grained and statistically sound overview of trial-by-trial changes in cue utilization. The increased precision offered by the DLM in tracking participants' cue utilization is important, as differences in task structure can have subtle effects on learning, which might otherwise be missed.

A graphical representation of the DLM is given in the left hand side of Figure 2. Like the classic lens model, the DLM consists of two regression models, one regressing the criterion and one regressing the responses onto the cues. The main difference is that we explicitly model the process by which the regression coefficients (the validity and utilization weights) change over time. This is indicated in Figure 2 by the arrows from the current cue validity and utilization weights to those at the next time point. There are only arrows between consecutive time points, which reflects the so-called first-order Markov assumption that current valid-

ity/utilization weights depend only on their previous values. Besides this assumption regarding the “memory” of the process, we do not assume there is a consistent trend in the changes; at each time point, the changes follow a Normal distribution with a zero mean. Such a process is also called a random walk.

To estimate the cue utilization weights, we use the Kalman technique introduced earlier. To be more precise, we actually use what is called the Kalman smoother, which gives more precise estimates than the Kalman filter. While the Kalman filter provides estimates of cue utilization using only current and earlier data, the Kalman smoother also takes later data into account. These later data are informative, because knowing what happened next can tell us something about the likelihood of an event. The Kalman smoother is used in an “off-line” setting, when estimation can be deferred until all data is collected. “On-line” estimation, when an estimate is needed as soon as a new datum is collected, must rely on the Kalman filter.

Estimation of cue utilization by the Kalman filter and smoother requires that certain parameters of the DLM are known. The first two parameters are the mean vector and variance-covariance matrix of the prior distribution of utilization weights. These parameters reflect a participant’s initial guess of the effect of the cues on the criterion, and the uncertainty associated with this guess. The third parameter is the variance-covariance matrix of the changes in cue utilization. The variances in this matrix reflect the average magnitude of trial-by-trial changes in cue utilization, while the covariances reflect how changes for one cue are correlated with changes for another cue. The final parameter is the variance of responses around their expectations. This parameter reflects how consistently participants respond according to their prediction strategy (as defined in terms of the cue utilization weights). For more details on the estimation of these parameters, and the Kalman filter and smoother technique, the reader is referred to Appendix A.

We have only mentioned estimation of the cue utilization weights here. While similar techniques can be used to estimate the cue validity weights, this was not necessary in our experiments, as the exact values of the cue validity weights were determined beforehand. However, a Kalman filter for cue validity estimation will be introduced later when we consider formal models of participants’ learning in MCPL tasks.

Types of change

Changes can be gradual and smooth, or abrupt and discontinuous. Previous research has mainly focused on the latter type. For example, there is an extensive literature in associative learning which draws upon experiments in which stimulus-reward relations change abruptly. In addition, concept learning paradigms such as the Wisconsin Card Sorting Test (Grant & Berg, 1948; Milner, 1963) involve category structures which change drastically from one trial to the next. The preference for abrupt changes in psychological research does not appear to be based on ecological concerns. Many processes in our environment are subject to gradual change. For instance, mortgage rates change gradually, as do cultural lore and the quality of vision. At times, however, these may undergo abrupt changes, as after a market crash, an exotic vacation, or putting on swimming goggles. Thus, people will encounter a mix of gradual and abrupt changes. The question then is whether different types of change are adapted to differently.

One possibility with abrupt changes is that they result in switching between representations (e.g., Bouton, 2004; Elliott & Anderson, 1995; Redish, Jensen, Johnson, & Kurth-Nelson, 2007). Rather than adapting a representation of the environment to inferred changes in the environmental structure, the previous representation is abandoned altogether in favour of a new

one. Bouton (2004) argued for such representational switching in the context of extinction studies. In these studies, an animal is first trained in a context in which a stimulus is paired with a reward. Then it is placed in a different context, where the stimulus is no longer paired with reward. As a result, the animal learns to suppress its response to the stimulus. When the animal later returns to the first context, it will renew its response to the stimulus (without needing to completely relearn this association). As such, there is evidence that the animal did not “unlearn” the association between the stimulus and response, but rather learned that this association was different in different contexts.

In typical extinction studies, the change is not only abrupt, but also relatively large, constituting a change from consistent reinforcement to a consistent absence of reward. Such deterministic reward schedules increase the salience of the change. As is well known, when extinction follows a partial reinforcement schedule, responses adapt less quickly to the new situation (e.g., Rescorla, 1999). This effect is understandable if learning is driven by surprise, as a number of consecutive non-reinforced trials is less surprising given a history of partial reinforcement than given a history of continuous reinforcement. Yu and Dayan (2005) distinguish between expected and unexpected uncertainty. The first type is associated with cue validity, and reflects the inherent uncertainty when basing inference on probabilistic cues. The second type is the uncertainty due to unexpected structural changes in cue validity. Abrupt changes will result in a relatively large unexpected uncertainty, while gradual changes will result in relatively little unexpected uncertainty. If the level of surprise (unexpected uncertainty) determines learning, we should expect slower learning and less adaptation in a gradually changing environment. We must be careful not to confound the rate of adaptation with the required adaptation. But we can compare adaptation in a condition in which, in a specific interval, the environment undergoes an abrupt change from one state to another, to adaptation in a condition in which, during an identical interval, the environment undergoes the same overall change, but in a more gradual fashion. If gradual changes result in a smaller rate of adaptation irrespective of the required change, we should expect participants in the gradual change condition to be less adapted at the end of the interval. This hypothesis will be investigated in the first two experiments presented here.

Insight

Intriguingly, even though people can master MCPL tasks quite well, it is often argued that they do so without insight into what they learned (Evans, Clibbens, Cattani, Harris, & Dennis, 2003; Gluck, Shohamy, & Myers, 2002; Harries, Evans, & Dennis, 2000; Nisbett & Wilson, 1977; Slovic & Lichtensteln, 1971; Wigton, 1996; York, Doherty, & Kamouri, 1987). Consequently, it has been argued that these tasks are learned implicitly (e.g., Berry & Broadbent, 1984; DeShon & Alexander, 1996; Knowlton, Mangels, & Squire, 1996; Seger, 1994). That is, participants learn to use the cues according to their validity without conscious access to what they have learned. However, previous studies have often used insensitive and invalid measures to assess explicit knowledge (Lagnado et al., 2006; Lovibond & Shanks, 2002). When probed in more valid ways, participants actually show good insight into the task structure (Lagnado et al., 2006; Newell, Lagnado, & Shanks, 2007; Speekenbrink, Channon, & Shanks, 2008). Some have suggested that, although participants may develop insight into their learning, they will do so only at a relatively late stage. If this is true, then what are the possibilities for insight when the environment changes? Will it necessarily lag behind? In the present study, we will use sensitive measures to assess participants' insight into the changing structure of the environment.

Overview

In the remainder of this article, we first present two experiments on human multiple-cue learning in a changing environment. In the first experiment, we compared participants' performance in an environment with abrupt changes to that in an environment with gradual changes, to assess whether participants adapted differently to these changes. In the second experiment, in addition to this, both the abrupt and gradually changing environments returned suddenly to a previous state. This was done to assess whether learning in a gradually changing environment would impede adapting to later abrupt changes. We then discuss five formal learning models and compare their fit to the data. A main distinction between these models is whether they implement global or local forms of learning. After discussing the fit of these models to the data of the first two experiments, we present results from a third experiment, which was designed especially to discriminate between global and local learning models.

Experiment 1

Method

Participants and apparatus. Twenty-eight participants took part in the experiment. All participants were recruited from the University College London subject pool and paid a basic fee (6 pounds) for their time, as well as an additional fee (between 0 and 5 pounds) based on their performance. Participants were randomly assigned to either the abrupt or gradual change condition. The two conditions did not differ in the distribution of sex (abrupt: 7 male, 7 female; gradual: 8 male, 6 female; $\chi^2(1) = 0, p = 1$), or mean age (abrupt: $M = 23.86, SD = 10.91$; gradual $M = 28.86, SD = 12.97$; $t(26) = 1.10, p = .28$). The entire experiment was run on a PC compatible desktop computer, using custom software written in C# for the .NET 2.0 framework. Participants were tested individually in a sound-dampened room.

Task. In the stockmarket prediction task, participants were asked to predict changes in the share price of a fictional software company "Mango" (the criterion) from changes in the share price of two related companies, "iTech" and "jSoft" (the cues). The predictive relation between the cues and criterion was as specified in the following equation:

$$Y_t = v_{1t} \times x_{1t} + v_2 \times x_{2t} + e_t \quad (1)$$

where Y_t denotes the value of the criterion on trial t (i.e., the change in Mango's share price from trial $t - 1$ to trial t), x_{1t} and x_{2t} denote the values of the cues (i.e., the changes in share price of iTech and jSoft), and e_t is a Normally and independently distributed error term with a mean of 0 and standard deviation of 3. The cue validity coefficients v_{1t} and v_2 determine the magnitude of the effect of the cues on the criterion. Throughout the task, the validity of the second cue was fixed at $v_2 = .4$, such that for example a 50 point increase in share price of this cue would increase the share price of Mango by 20 points. The validity of the first cue varied over trials t . In both conditions, the validity of the first cue started with a negative value of -0.8, so that for example a 50 point increase in share price of this cue would *decrease* the share price of Mango by 40 points. During the task, the validity changed from its initial value of -0.8, via 0, to a final value of 0.8, in discrete steps at trials 101 and 201 in the abrupt change condition, and linearly in the gradual change condition (see Figure 5 for a graphical depiction of the cue validity patterns).

More formally, the cue validity pattern in the abrupt change condition can be expressed as in the following equation:

$$v_{1t} = \begin{cases} -0.8 & \text{for } t = 1, \dots, 100 \\ 0 & \text{for } t = 101, \dots, 200 \\ 0.8 & \text{for } t = 201, \dots, 300 \end{cases}$$

In the the gradual change condition, the equation for the cue validity is

$$v_{1t} = v_{1(t-1)} + 1.6/300$$

starting at the initial trial $t = 1$ with a value of $v_{1(1)} = -0.8$.

The values of the cues, x_{1t} and x_{2t} , were sampled independently from uniform distributions over the range from -100 to 100. The criterion, Y_t , was generated by inserting the cue values in Equation 1, with the value of the error, e_t , drawn from a Normal distribution with mean 0 and standard deviation 3. Pairs of participants in the abrupt and gradual change condition were yoked so that they received the same sequence of cue values and error values.

Procedure. At the start of the experiment, participants were given on-screen instructions that they were to take the role of a financial analyst for an investment company. They had to predict the price of shares in software company Mango and the investment company would gain or lose money based on the accuracy of their predictions. Participants were told they could base their predictions on information about the share price of the related companies iTech and jSoft. They were also informed that it would be difficult to predict the share price with complete accuracy, but that their predictions should improve with practice. Participants were not informed that the relation between a cue and the criterion would change over time.

The task consisted of 300 trials. On each trial, participants observed the values of the two cues both numerically and in graphical form as filled bars (see Figure 3). They could predict the change in share price of Mango by moving a slider. Participants could view the effect of this change on the new share price in a graph, which also showed the share price on the previous five trials. In addition, the predicted change in share price was shown graphically as a filled bar. After recording their prediction, participants received feedback as to the actual share price. The change in share price was shown in the bar beside the slider and the resulting share price was shown in the graph. The difference between the actual and predicted change in share price was indicated by a coloured block in the bar on the left hand of the slider. Additional feedback was given in a message in the “console”, which contained the numerical values of the predicted share price, the true share price, the difference between these values, and the amount won or lost by the investment company. The earnings for the investment company were determined by a linear function of the absolute deviation between a participant’s prediction and the actual value of the criterion:

$$\text{earnings}_t = 100 - 5 \times |R_t - Y_t|,$$

where R_t denotes a participant’s response, and Y_t the actual value of the criterion. Thus, in order to make money for the investment company, the difference between participants’ predictions and the actual value of the criterion should be less than 20. The company started with a total fund of 20,000, and the earnings on each trial were added or subtracted from the total funds. Participants could keep track of the current funds by means of a bar on the upper right of the screen.

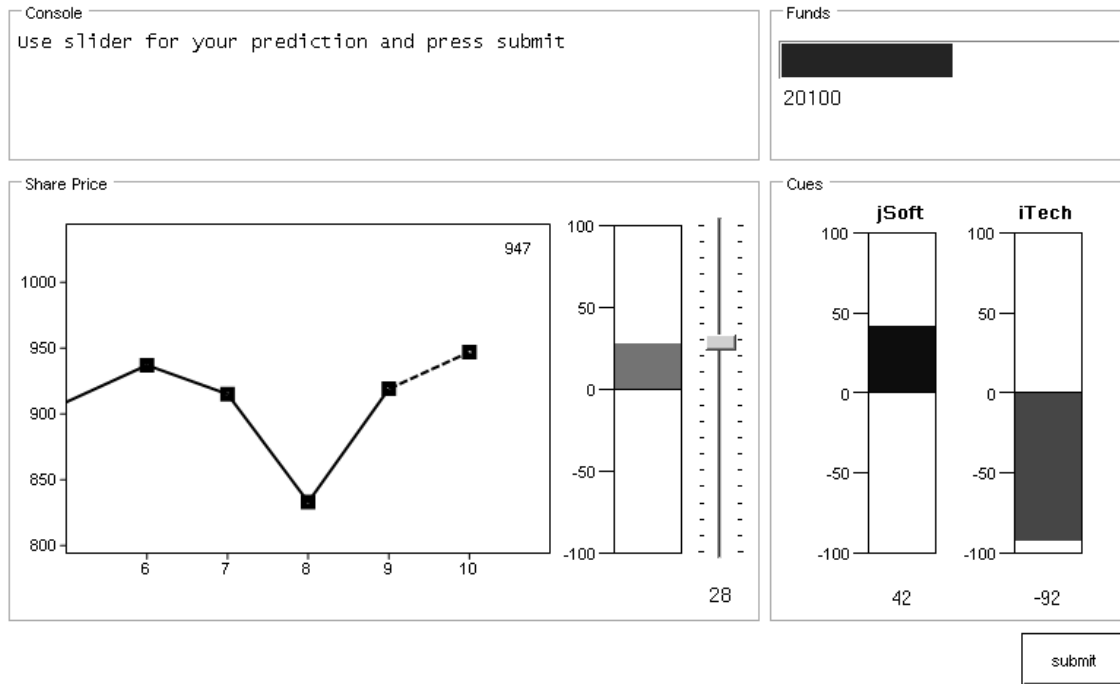


Figure 3. Representation of the screen display in the Stockmarket Prediction Task.

The task was divided into four blocks of 75 trials each. At the end of each block, participants were asked to rate the strength of the relation between the cues and criterion on a scale from -10 (highly negative) to 10 (highly positive), and asked to choose which of the two cues had the strongest effect on the criterion. To avoid cueing participants to the change in one of the cue-criterion relations, the start of each block (trials 1, 76, 151, 226) was deliberately asynchronous with the change in cue validity in the abrupt change condition (after trials 100 and 200). The order of the two cues on the screen, and the cue labels, were randomized for each participant.

Results and discussion

The structure of this section is as follows. First, we will assess participants' learning performance and test for differences between the conditions. We will then focus on cue utilization more directly. Finally, we will investigate participants' insight into the changing task structure.

Performance. Individual correlations between predictions and criterion values ranged from 0.04 to 0.92 ($M = 0.65$, $SD = 0.26$) in the abrupt change condition, and from 0.42 to 0.83 ($M = 0.59$, $SD = 0.12$) in the gradual change condition. There was one participant (A6) in the abrupt change condition for whom the correlation was not reliably different from 0, $t(298) = 0.66$, $p = .51$ (correlations were significantly larger than 0 for all other participants). All further analyses were performed both with and without inclusion of participant A6. The results of the full analyses are reported; exclusion of participant A6 did not change the qualitative

pattern of results.

We used the Mean Absolute Deviation (MAD) of participants' predictions from the actual criterion values as the main measure of performance. A 2 (condition) \times 15 (block: 20 trials each) ANOVA with repeated measures on the second factor showed, besides a significant effect of block, $F(14, 364) = 24.44$, $MS_e = 72.6$, $p < .001$, a significant interaction between condition and block, $F(14, 364) = 2.73$, $p < .001$. The main effect of block reflects a general decrease in error from the first block ($M = 39.28$, $SD = 9.56$) to block 10 ($M = 12.83$, $SD = 7.51$), after which it slightly increases until the last block ($M = 20.01$, $SD = 11.57$). In the abrupt condition, there were sudden increases after the changes in cue validity, while the pattern was more smooth in the gradual change condition. The main effect of condition was not significant (abrupt: $M = 22.81$, $SD = 24.48$; gradual: $M = 21.96$, $SD = 20.64$), $F(1, 26) = 0.14$, $MS_e = 525.2$, $p = .71$. Thus, while the difference between the conditions in cue validity pattern resulted in a difference in the pattern of error over trials, average performance was equal in the two conditions.

Which cues were used? Participants may have based their predictions on more information than the two relevant cues (the changes in share price of jSoft and iTech). To increase the realism of the task, participants were shown a graph of Mango's share price on the five preceding trading days, and participants may have been influenced by this information. Participants may also have assumed that the effect of one cue was moderated by the other cue (i.e., a cue interaction). To investigate these possibilities, we fitted four different DLMS: the Correct (C) model, in which only the two relevant cues were included, the Previous change (P) model, which included the previous change in share price as an additional cue, the Interaction (I) model, which included the cue interaction as an additional cue, and the Interaction and Previous change (I+P) model, which included both additional cues. An additional concern in this analysis was the level of individual differences. Do participants adapt their cue utilization at the same rate? Does this rate depend on whether the changes in cue validity were gradual or abrupt? Do participants apply their prediction strategies with the same consistency? To answer these questions, we fitted the four models to different subsets of the data, thereby fixing the model parameters to be identical for all participants, identical within conditions, or allowing different parameter values for each participant.

Full details of the analysis are given in Appendix B. In summary, the results indicated considerable individual variation in the model parameters, highlighting the need for individual based analysis of cue utilization. Overall, the large majority of participants ($n = 23$) relied on only the two relevant cues, while some participants ($n = 4$) based their predictions also on the previous change in share price or assumed a cue interaction ($n = 1$). Inspection of the utilization of the previous change in share price showed that it had a positive effect for most participants who relied on it, predicting an increase (decrease) in share price if the share price previously increased (decreased). These participants seem to have assumed a form of "momentum" in the share price process. However, utilization of the previous change in share price tended to decrease during the task, and most participants who initially used this information learned to ignore it (the utilization weight approached 0 by the end of the task).

Cue utilization. Estimates of participants' cue utilization were obtained from the best fitting DLMS identified above. The resulting patterns of cue utilization over trials, or cue utilization profiles, are depicted in Figure 4. In accordance with the previous analysis, there was considerable individual variability in the utilization profiles. For the majority of participants, cue utiliza-

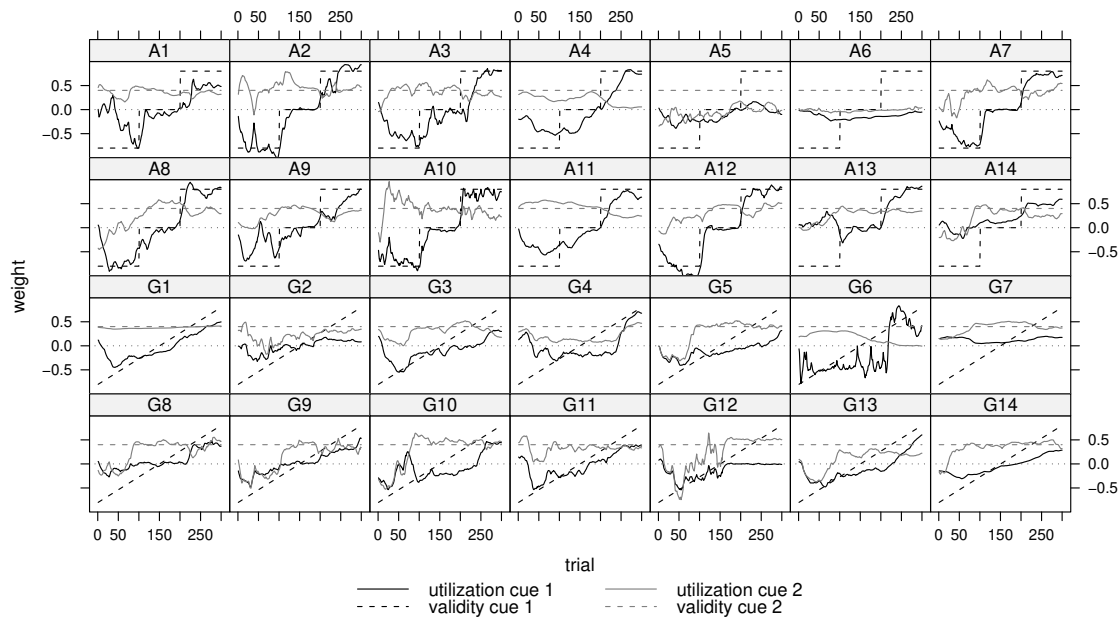


Figure 4. Cue utilization by participants in the abrupt change condition (A1-A14) and gradual change condition (G1-G14). Utilization (solid lines) was estimated with model C, except for participants A1, A6, A12 and G7 (model P) and participant G14 (model I). Validity (broken lines) represents cue validity.

tion was strongly related to the changes in cue validity, although some (e.g., A5 and A6) made relatively little use of the cues. Especially for some participants in the abrupt change condition (e.g., A7, A8, A10 and A12), the utilization of cue 1 closely followed the “step function” of the cue validity. The utilization of cue 2 was variable as well, but, after an initial stage, appeared to vary around a horizontal line often located at cue validity. Interestingly, some participants (A4 and G8) initially learned to use this cue in accordance with its validity, but stopped utilizing it towards the end of the task. This decrease in utilization appears to be accompanied by an increase in the utilization of cue 1, which is reminiscent of the cue competition effect found by Bussemeyer, Myung, and McDaniel (1993a). They showed that an increase in the (absolute) validity of one cue resulted in a decrease of the (absolute) utilization of a second cue. This competition between cues is consistent with the blocking effects found in associative learning experiments. In the present experiment, cue competition seems limited to a few participants.

The average utilization curves are depicted in Figure 5. A condition \times block (30 blocks of 10 trials each) ANOVA was conducted on the utilization of cue 1, with repeated measures on the last factor and a Greenhouse-Geisser correction for heterogeneity of variance. A main effect of block, $F(2.41, 62.55) = 65.30$, $MS_e = 0.22$, $p < .001$, confirms that participants adapted their utilization of cue 1 during the task. In addition, a condition \times block interaction shows that this pattern depended on whether the changes in cue validity were abrupt or gradual, $F(2.41, 62.55) = 8.90$, $p < .001$. A similar ANOVA showed that participants also adapted their utilization of cue 2 during the task, $F(2.57, 66.85) = 12.71$, $MS_e = 0.12$, $p < .001$. There was a trend towards an interac-

tion between condition and block, $F(2.57, 66.85) = 2.51$, $p = .07$, indicating that the form of the change in cue 1 validity may have affected utilization of cue 2. Figure 5 shows that the difference between the conditions is located at the start and end of the task. In the abrupt change condition, the utilization of cue 2 shows a gradual rise towards cue validity. In the gradual change condition, utilization appears to linger around 0 until trial 50, followed by a more rapid rise towards cue validity. Interestingly, after this initial stage, the utilization appears to remain stable for the gradual change condition, while the average utilization in the abrupt change condition drops somewhat after trial 200 (corresponding to the second change in the validity of cue 1).

Both Figure 5 and the results of the ANOVA suggest that participants' utilization adapted to the specific cue validity profile of their condition. Cue utilization appeared closer to a step function in the abrupt change condition, and more like a linear function in the gradual change condition. To test this more directly, we estimated two regression models for the utilization of cue 1, one with the cue validity profile of the abrupt change condition, and the other with cue validity profile of the gradual change condition as predictor. The regression models were fitted to the data of individual participants. For each condition, we then compared the (absolute) residuals of the models. Adaptation to the form of change should result in smaller residuals when using the relevant cue validity profile (i.e., the abrupt profile for the abrupt change condition, and the gradual one for the gradual change condition). Paired sample Wilcoxon tests confirmed that the abrupt validity curve described utilization better than the gradual change curve in the abrupt change condition, $Z = -12.33$, $p < .001$, while the gradual validity profile fitted better in the gradual change condition, $Z = 10.84$, $p < .001$.

Figure 5 indicates that utilization of cue 1 at the end of the task was stronger in the abrupt than gradual change condition. A t-test comparing the utilization at trial 300 between the conditions confirmed this, $t(26) = 2.69$, $p < .05$. The utilization of cue 2 did not differ between the conditions at this point, $t(26) = t = -1.09$, $p = 0.29$. Hence, utilization was closer to cue validity when the latter changed abruptly than when it changed gradually. This is consistent with a smaller rate of adaptation following gradual change.

Insight. Participants' insight into the task structure was assessed with two kinds of questions. In the rating questions, participants were asked to rate the strength and direction of each cue's effect on the criterion. Figure 6 shows the mean ratings together with the (rescaled) validity weights at the time of rating. The average ratings follow the cue validity patterns quite well; apart from trial 225 in the gradual condition, the average ratings were in the expected directions, and in most cases significantly so. The relation between ratings and cue validity was confirmed by a regression analysis without intercept¹, which showed a significant effect of cue validity, $F(1, 223) = 110.90$, $p < .001$, $R^2 = 0.33$. Interestingly, a similar regression model with cue utilization as predictor showed that this had a stronger effect on the ratings, $F(1, 223) = 256.3$, $p < .001$, $R^2 = 0.53$. In fact, while the addition of cue utilization to the model with cue validity resulted in a significant reduction of error, $F(1, 222) = 97.12$, $MS_e = 11.42$, $p < .001$, adding cue validity to the model with cue utilization as predictor did not, $F(1, 222) = 0.32$, $p = .57$. As such, participant's ratings seem to reflect their inference of the task structure, rather than the task structure directly.

¹We did not include an intercept in the model because we were interested in whether participants could detect the direction of the cue-criterion effects as well as their strength. Including an intercept would have resulted in a model in which positive ratings do not necessarily reflect positive validity weights.

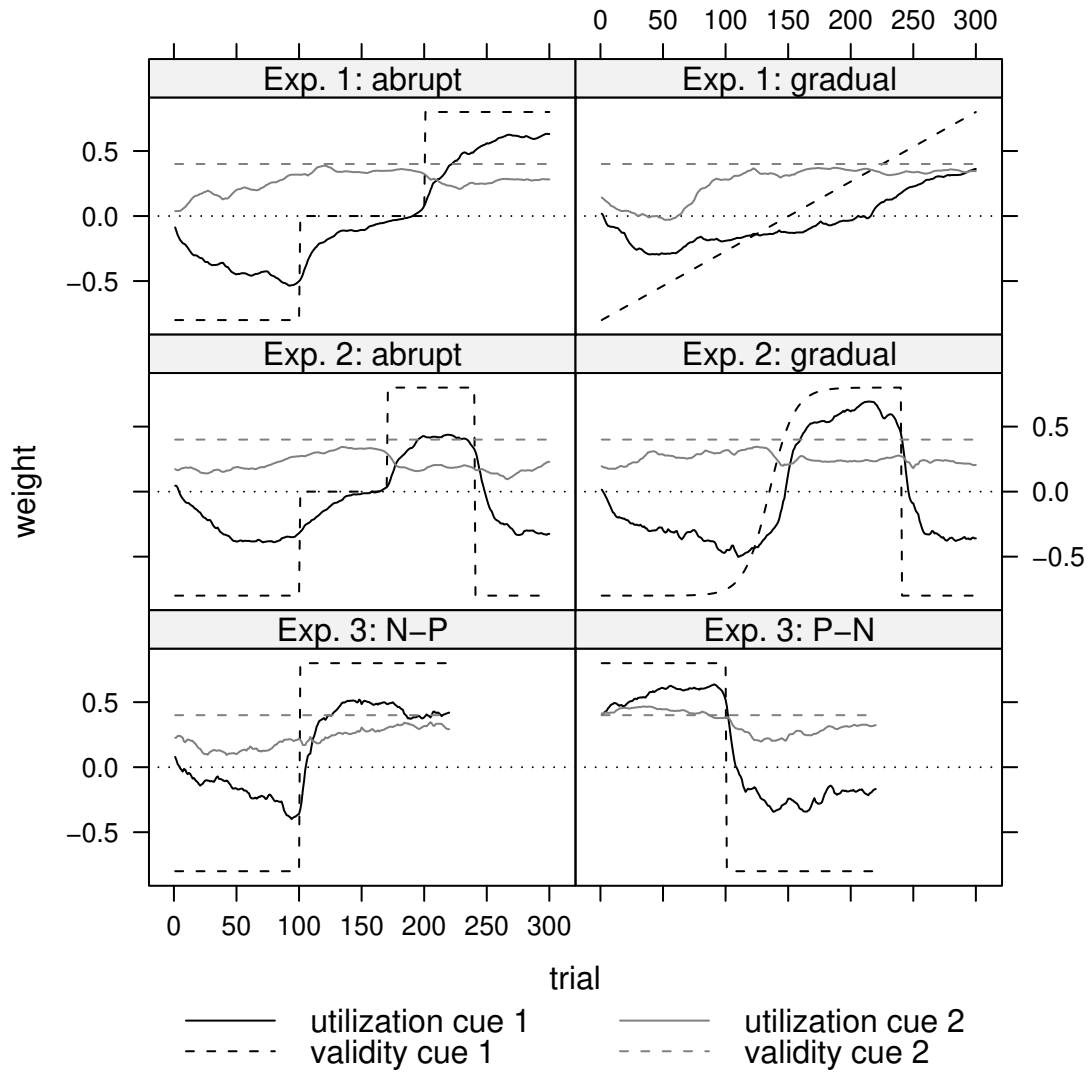


Figure 5. Average estimated cue utilization and actual cue validity.

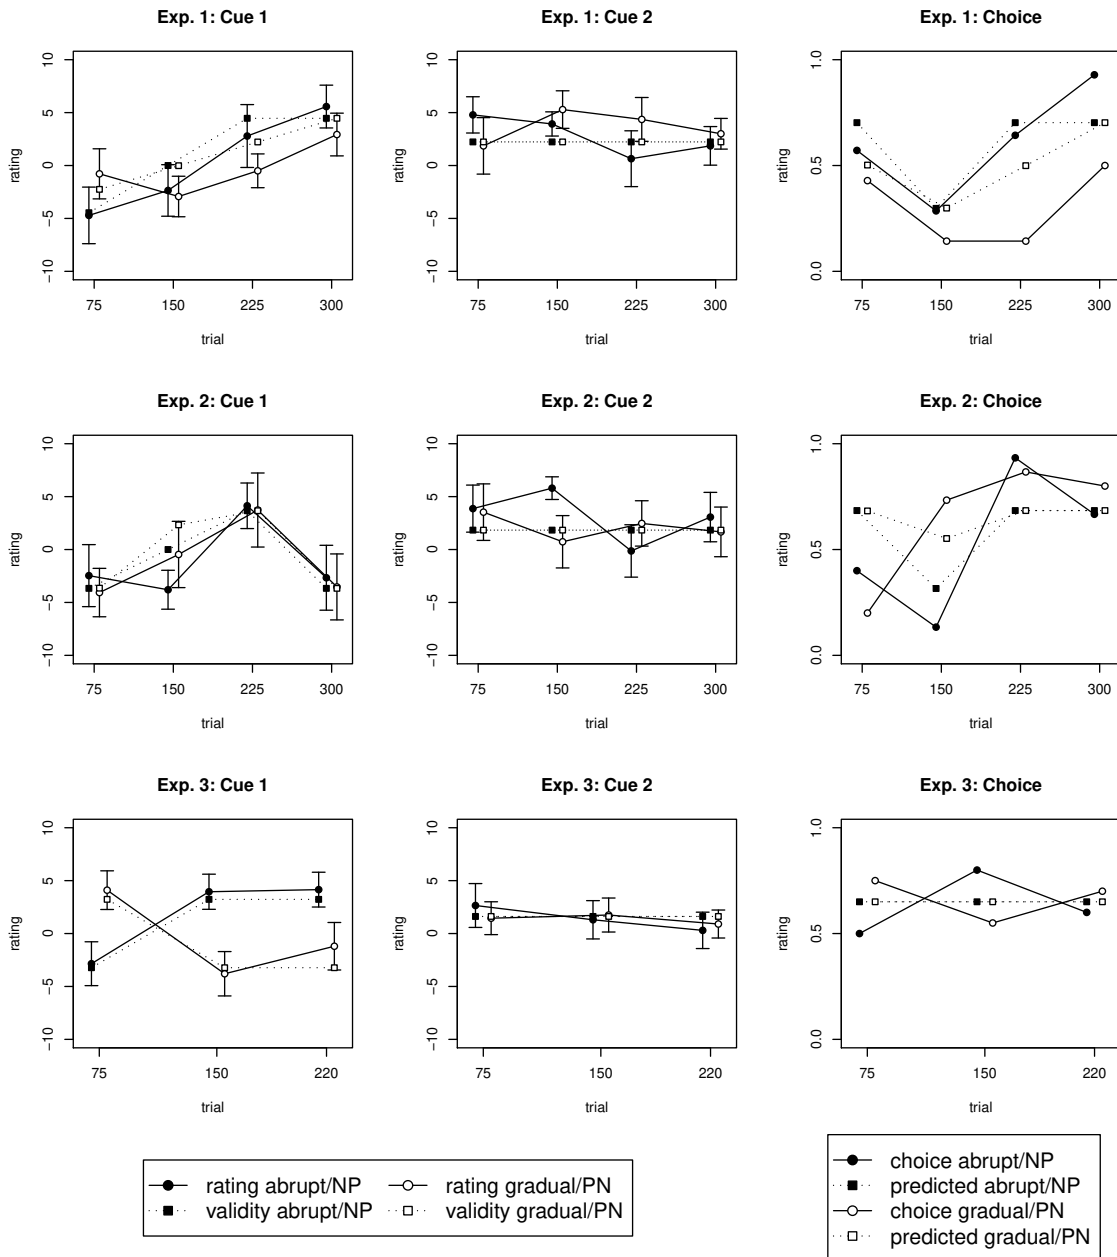


Figure 6. Means and 95% confidence intervals of the strength ratings and validities for the two cues in Experiments 1–3. Cue validities were rescaled by linear regression without intercept. To allow for easier discrimination, the ratings (and validities) in the two conditions have been slightly displaced on the x-axis and the ratings and validities at the different trials are connected by lines. The reader should note that these lines are not indicative of the actual cue validity profiles (e.g., the validity of cue 1 in the abrupt change condition changed abruptly on trial 100, see Figure 5).

In the forced choice questions, participants were asked to pick the cue with the strongest effect on the criterion. To assess the extent to which the choices reflect actual differences in cue validity, we fitted a logistic regression model with choice (1 = cue 1, 0 = cue 2) as dependent variable and the difference in absolute validity ($d = |v_{1,t}| - |v_{2,t}|$) as predictor (there was no intercept). This showed a significant effect of the difference in validity, $b = 2.15$ (S.E. = 0.60), $p < .001$, although the model did not fit the data very well, $\chi^2(111) = 141.11$, $p < .05$. We also fitted a similar model with the difference in absolute utilization (rather than validity) as predictor. This model did fit the data, $\chi^2(111) = 125.76$, $p = .16$, and showed a significant effect of utilization, $b = 4.08$ (S.E. = 0.88), $p < .001$. As for the ratings questions, while the addition of the cue utilization difference to the model with cue validity difference significantly increased model fit, $\chi^2(1) = 16.23$, $p < .001$, the reverse was not true, $\chi^2(1) = 0.87$, $p = .35$. Again, these results indicate that participants' insight was based on their inference of the task structure, rather than the task structure directly.

In summary, we found no evidence for a dissociation between participants' knowledge of the environment as manifest in their predictions, and their explicit judgements of the effects and relative importance of the cues. As such, we found no support for implicit learning in this task.

Experiment 2

Experiment 1 showed that participants were responsive to the changes in cue validity. Sudden and gradual changes in cue validity resulted in correspondingly shaped patterns of cue utilization. While there was no overall performance difference between the conditions, participants' utilization at the end of the task was closer to cue validity in the abrupt change than in the gradual change condition. This is consistent with slower adaptation to gradual changes, as the overall change in cue validity was identical in both conditions. However, the conditions differed in another respect beside the type of change. In the abrupt change condition, cue validity was initially stable, while cue validity changed from the outset in the gradual change condition. As participants had no prior knowledge of the task structure, they needed to form a representation of the structure whilst simultaneously adapting it to the changes. In contrast, participants in the abrupt change condition could form a stable representation before needing to adapt to a change in cue validity. As it is likely that change is more easily detected from a relatively stable representation of the task, participants in the abrupt change condition may have been advantaged. The second experiment addresses this issue by starting the task with a period of relative stability in both conditions. In addition, the end of the task included an identical abrupt change in both conditions, in which the validity of the changing cue returned to its original state. Insofar as abrupt changes result in the formation of different representations of the task, rather than the modification of a single representation (e.g. Bouton, 2004; Elliott & Anderson, 1995; Redish et al., 2007), we should expect participants in the abrupt change condition to be advantaged in adapting to this final change, as switching between representations should be quicker than adapting a single representation.

Method

Participants. Thirty students from University College London took part in the experiment. Participation was voluntary, and participants were entered into a draw to win one of two cash prizes (20 and 10 pounds). Participants were randomly assigned to either the abrupt or gradual

change condition. The two conditions did not differ in the distribution of sex (abrupt: 5 male, 10 female; gradual: 4 male, 11 female; $\chi^2(1) = 0, p = 1$), or mean age (abrupt: $M = 21.20, SD = 2.08$; gradual $M = 20.53, SD = 1.19$; $t(28) = 1.08, p = .29$).

Materials. The task was similar to that of Experiment 1. The only difference was in the validity of cue 1. In the abrupt change condition, the cue validity profile was similar to Experiment 1, starting at a negative value of $v_{1t} = -0.8$ (trials 1 to 100), then changing to a value of $v_{1t} = 0$ (trials 101 to 170), and then to a positive value of $v_{1t} = 0.8$ (trials 171 to 240). However, in this experiment, the validity then returned to its initial negative value of $v_{1t} = -0.8$ (trials 241 to 300). In the gradual change condition, the validity of cue 1 also went from a negative value of $v_{1t} = -0.8$ to a positive value of $v_{1t} = 0.8$, but this change was smooth, slowly changing at the start and end, but faster in the middle. The end of the task included the same abrupt change from $v_{1t} = 0.8$ to $v_{1t} = -0.8$ as in the abrupt change condition. More formally, the cue validity profile in the gradual change condition can be expressed in the following equation:

$$v_{1t} = \begin{cases} -0.8 + \frac{1.6}{1 + \exp(14.3 - t/10)} & \text{for } t = 1, \dots, 240 \\ -0.8 & \text{for } t = 241, \dots, 300 \end{cases}$$

The logistic function which describes v_{1t} between trials 1 and 240 has the appearance of an S-shaped curve: relatively stable at first, then changing increasingly rapidly, after which the rate of change decreases again to a period of relative stability. Thus, while continuously changing, the validity starts and ends with a period of relative stability, just as in the abrupt change condition. With this particular function, 95% of the change occurs in the space of 77 trials (from trial 95 until trial 172). A graphical representation of the resulting cue validity patterns can be found in Figure 5.

Procedure. The procedure was identical to that of Experiment 1, apart from a section in the instructions, in which participants were informed that they would enter a draw to win cash prizes, where the chance of winning was based on their performance.

Results and discussion

Performance. The correlation between predictions and criterion values ranged from 0.01 to 0.81 ($M = 0.50, SD = 0.25$) in the abrupt change condition, and from -0.02 to 0.91 ($M = 0.57, SD = 0.25$) in the gradual change condition. There was one participant in each condition (A13 and G11) for whom the correlation was not significantly different from 0, $t(298) = -0.39, p = .70$ and $t(298) = 0.12, p = 0.91$, respectively. The correlation was significant for all other participants. As before, we report the results for all participants; unless stated otherwise, the same pattern of results was obtained from the analyses after exclusion of A13 and G11.

The MAD between predicted and actual criterion values was analysed with a 2 (condition) \times 15 (block: 20 trials each) ANOVA with repeated measures on the second factor and a Greenhouse-Geisser correction. This analysis showed a significant effect of block, $F(5.86, 164.08) = 14.76, MS_e = 2.39, p < .001$, and a significant interaction between condition and block, $F(5.86, 164.08) = 2.93, p < .05$. The main effect of condition was not significant (abrupt: $M = 31.32, SD = 31.05$; gradual $M = 30.93, SD = 31.24$), $F(1, 28) = 0.01, MS_e = 116.57, p = .91$; hence, as in Experiment 1, the form of the validity changes did not affect overall performance, but did affect the pattern of performance over the task.

Which cues were used? The results replicated those in Experiment 1 (see Appendix B for details). The large majority of participants ($n = 24$) relied on just the two cues when making predictions, while some participants ($n = 4$) also relied on the previous change in share price or assumed an interaction between the cues ($n = 2$). In addition, there was strong evidence for individual variation in the model parameters.

Cue utilization. The estimated cue utilization profiles are given in Figure 7. As in Experiment 1, participants' cue utilization adapted to the changing environment. Again, there was a lot of individual variability; some participants (e.g., A8 and G9) adapted quickly and accurately to the changes, while others (e.g., A10 and G12) showed relatively little adaptation. Some participants (e.g., A11, G1, G5, G11 and G14) appeared to have initial difficulties in utilizing cue 1 according to its (negative) validity. This is consistent with the often noted difficulty in learning negative cue-criterion relations (e.g., Brehmer, 1973; Naylor & Clark, 1968). Participants G5, G11 and G14 show a sudden large change towards negative utilization, before changing towards positive utilization (in accordance with the change in validity). After initially neglecting this cue, these participants seemed to suddenly realise the negative relation between this cue and the criterion. A number of participants (A2, A5, A6, A9, G1, G5, G8 and G10) show evidence of cue competition, in that the utilization of cue 2 clearly drops in the period where cue 1 is at its maximum positive validity. Interestingly, cue competition seems strongest in the period in which both cues have a positive effect on the criterion; after the final change, when cue 1 is at its maximum negative validity, the utilization of cue 2 rises again. As such, cue competition may be restricted to situations in which cues have a similar effect on the criterion and differentiation between them is relatively difficult.

Figure 5 shows the average utilization in the two conditions. The shape of the average utilization curves appears to match the shape of the two validity curves, with two inflection points in the abrupt change condition corresponding to the abrupt changes in cue validity, and a more gradual change pattern in the gradual change condition. A condition \times block (15 blocks of 20 trials each) ANOVA for the utilization of cue 1, with repeated measures on the last factor and Greenhouse-Geisser correction, replicated the results of Experiment 1. There was a significant effect of block, $F(4.16, 116.58) = 64.98$, $MS_e = 0.18$, $p < .001$, as well as a significant interaction between block and condition, $F(4.16, 116.58) = 4.86$, $p < .001$. As before, differences in the cue validity patterns resulted in differences in cue utilization. The same ANOVA for the utilization of cue 2 showed a trend for block, $F(4.25, 119.07) = 2.25$, $MS_e = 0.09$, $p = .06$, and no significant effects.

The results of the regression analysis, predicting cue utilization from the cue validity profiles, also replicated those of Experiment 1. In the abrupt change condition, the abrupt change validity curve described utilization better than the gradual change curve, $Z = -8.19$, $p < .001$, while the reverse was true for the gradual change condition, $Z = 11.55$, $p < .001$.

We investigated the speed of adaptation by comparing utilization in the period before the final abrupt change. Figure 5 shows no evidence of a lower rate of adaptation in the gradual change condition. If anything, average utilization appears higher than in the abrupt change condition. Averaging cue utilization over the 20 trials preceding the final abrupt change, this difference was significant, $t(28) = 2.15$, $p < 0.05$. However, this result was due to the two sub-performing participants; after exclusion of A13 and G11, the difference was no longer significant, $t(26) = 1.87$, $p = .073$. There was thus no strong evidence that cue utilization was closer to

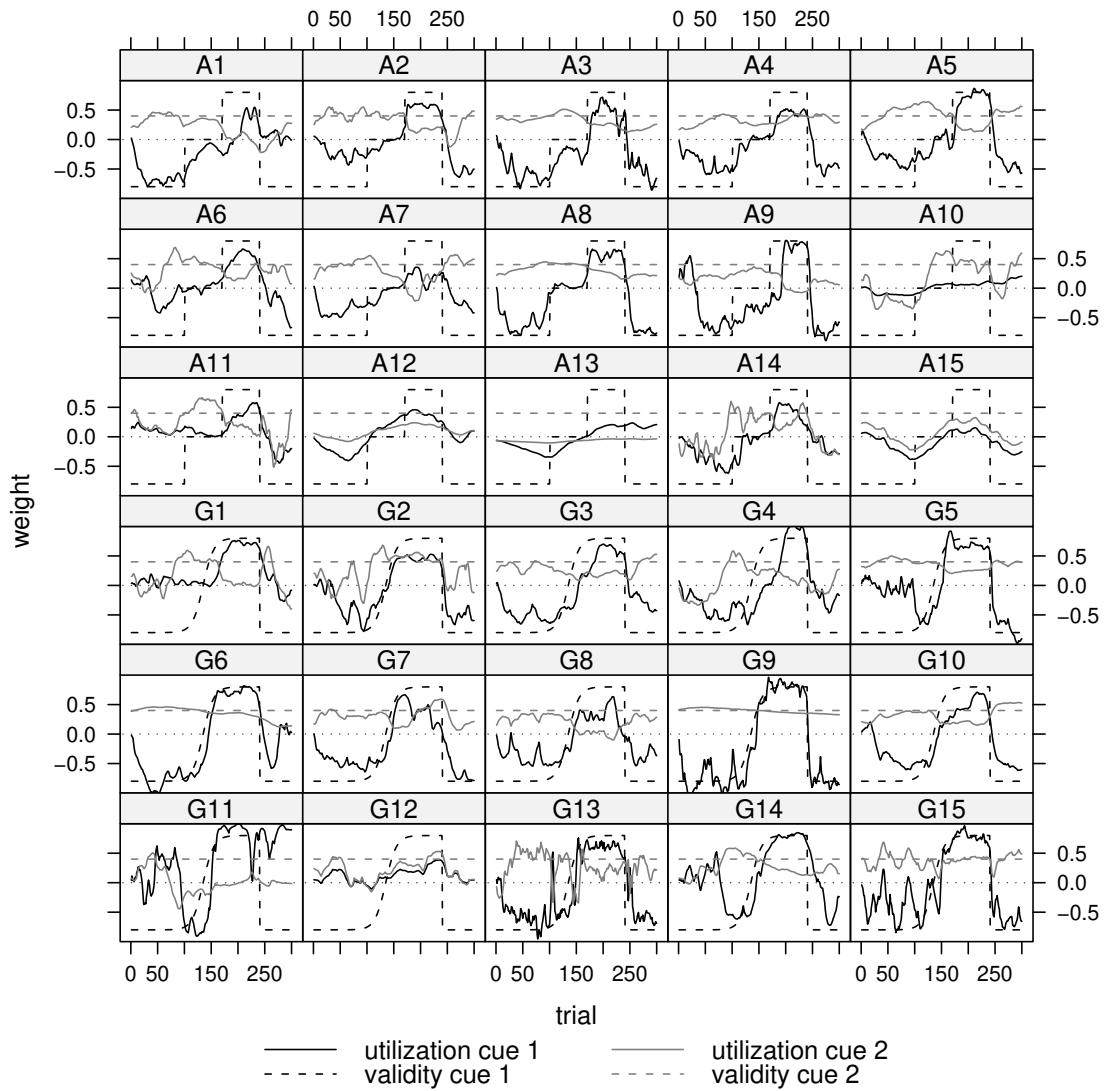


Figure 7. Cue utilization for participants in the abrupt change condition (A1-A15) and gradual change condition (G1-G15). Estimated utilization (solid lines) was computed with model C, apart from participants A6 and A7 (model I) and A6, G4, G5 and G12 (model P). Validity (broken lines) represents actual cue validity.

optimal in the gradual condition. In contrast to Experiment 1, we found no evidence of slower adaptation to gradual change.

To test the hypothesis that abrupt changes result in the formation of multiple representations of the task, we assessed adaptation to the final abrupt change in a condition by block (7 blocks of 10 trials each, starting at trial 230) ANOVA with repeated measures on the last factor and Greenhouse-Geisser correction. This showed only a significant effect of block, $F(1.87, 52.32) = 67.38$, $MS_e = 0.14$, $p < .001$. Neither the main effect of condition, $F(1, 28) = 0.02$, $MS_e = 0.53$, $p = .88$, nor its interaction with block, $F(1.87, 52.32) = 1.17$, $p = .31$, were significant. There is thus no evidence that the abrupt change condition adapted more quickly to the final abrupt change, as would be expected from reinstatement of a previous task representation.

Insight. The results of the insight questions are depicted in Figure 6. As before, the average ratings were mostly consistent with cue validity and the regression analyses replicated the results of Experiment 1. The effect of cue validity on the ratings was significant, $F(1, 239) = 63.54$, $p < .001$, $R^2 = 0.21$, but a regression model with cue utilization as predictor fitted better, $F(1, 239) = 128.5$, $p < .001$, $R^2 = 0.35$. And, while adding cue utilization to the model with cue validity as predictor resulted in a significant reduction of error, $F(1, 238) = 44.35$, $MS_e = 19.74$, $p < .001$, adding cue validity to the model with cue utilization as a predictor did not, $F(1, 238) = 1.45$, $p = .23$.

The results of choice questions also replicated those in Experiment 1. While showing a significant effect of validity, $b = 1.93$ (S.E. = 0.52), $p < .001$, the model with cue validity as predictor did not fit the data very well, $\chi^2(119) = 151.66$, $p < .05$. A model with cue utilization as predictor did fit the data, $\chi^2(119) = 141.67$, $p = .08$, and showed a significant effect of utilization, $b = 3.31$ (S.E. = 0.77), $p < .001$. And, while adding cue utilization to the model with cue validity significantly increased model fit, $\chi^2(1) = 12.22$, $p < .001$, the reverse was not true, $\chi^2(1) = 2.23$, $p = .14$.

In summary, the significant effects of validity and in particular utilization on participants' ratings and choices shows participants had a considerable degree of insight into what they learned. Again, we failed to find evidence for implicit learning in this task.

Learning models

The results of Experiment 1 and 2 show that participants were responsive to changes in the environment and learned to utilize the cues in accordance with their (changing) validity. By comparing different formal learning models, we will now investigate how they may have done so.

While there are many models of category learning, relatively few have been proposed to model how people learn continuous functions of multiple cues. We can distinguish two broad classes of function learning and categorization models: global and local learning models. Global learning models extract a single function (or rule) relating cues to the criterion that applies to all possible cue values. When this function does not make an accurate prediction for a cue pattern, the function is adjusted to reduce the prediction error. A defining characteristic of global learning models is that the function will be adjusted for the whole range of possible cue values. A change detected for for one cue value will thus generalize to all other cue values. Local learning models do not display such generalization. They learn localized associations between the criterion and specific cue patterns by strengthening associative links between cue patterns

and criterion values, or by storing and recollecting cue-criterion pairings as in exemplar models. In local learning models, new observations may change the predictions for the observed cue pattern, but importantly, there is little or no effect on predictions to other cue patterns.

We will now discuss different instantiations of global and local learning models suitable for multiple cue learning tasks with changing cue-criterion relations. We formulate two new global Bayesian models of MCPL, which differ in the assumed structure of the environment: the Bayesian Linear Filter, which assumes the cues combine additively in determining the criterion, and the Competitive Bayesian Filter, which assumes the cues compete to determine the criterion. As a final global learning model, we use a Least Mean Squares network, which makes similar assumptions as the Bayesian Linear Filter, but uses a different (non-Bayesian) form of learning. We compare the global learning models to two previously successful local learning models: the Generalized Context Model and the Associative Learning Model. The discussion will be mostly conceptual, describing the main ideas behind each model; technical details are given in Appendix C.

Global learning: Bayesian Linear Filter (BLF)

Bayesian modelling has become an increasingly popular framework to describe human inductive learning (Tenenbaum, Griffiths, & Kemp, 2006), reasoning (Oaksford & Chater, 2007), and causal inference (Griffiths & Tenenbaum, 2005), to name but a few areas. Bayesian models update probability distributions on the basis of incoming information and are well suited to online (e.g., trial-by-trial) learning. The Bayesian Linear Filter (BLF) is one possible Bayesian model. It assumes the environment can be modelled as in the DLM (the criterion is a linear function of the cues, and validity weights change over time according to a random walk). The BLF learns the cue validity weights by using the Kalman filter to compute the posterior distribution of cue validity weights. The Kalman filter has been previously proposed to model reward learning in conditioning (e.g., Dayan, Kakade, & Montague, 2000; Kakade & Dayan, 2002; Sutton, 1992). The BLF is related to this model, but learns the parameters of a linear function rather than the level of reward in the presence (versus absence) of a cue.

When making predictions, the BLF computes a predictive distribution of the criterion. The predictive distribution is based on the posterior cue validity distribution and more uncertainty about the cue validity weights results in more uncertainty regarding the criterion value. In addition, the variance of the predictive distribution increases with the extremity of the cue values; as predictions are the summed product of cue values and validity weights, cue validity uncertainty is magnified by higher the absolute cue values. We assume participants' responses are made in accordance with the predictive distribution.

Global learning: Competitive Bayesian Filter (CBF)

The BLF assumes the value of the criterion depends on all cues simultaneously. An alternative assumption is that the cues compete to determine the criterion; on each trial, only a single cue is responsible for the value of the criterion, but which cue that is can vary over time. These assumptions may underlie the cue competition effects (e.g., Birnbaum, 1976; Busmeyer et al., 1993a) referred to earlier. To learn in such a competitive environment, two problems must be solved: determine which cue is responsible for the criterion, and determine the effect of each individual cue on the criterion. These two tasks are performed optimally by what we call the

Competitive Bayesian Filter (CBF), which is closely related to the competitive mixture of experts model (Dayan & Long, 1998).

The CBF tracks changes in cue validity in a similar way to the BLF and uses posterior cue validity distributions to compute predictive distributions of the criterion. But in the CBF, each cue predicts a different value of the criterion. As there is uncertainty about which cue will be responsible on a given trial, the final predictive distribution is a mixture of the individual predictive distributions, weighting each by the probability that the corresponding cue will be responsible on that trial. After observing the actual criterion value, this information is used to compute a posterior probability of cue responsibility. The extent to which the CBF learns about the individual cue-criterion relations depends on this assessment of responsibility, as the observed value only provides information about the relation between the criterion and the responsible cue.

We assume participants' predictions are made by weighting the predictions for each individual cue according to their predicted responsibility. As a result, cues are relatively under-utilized when compared to their inferred cue validity. Furthermore, in situations where the assumption of cue competition is wrong, such as in the present task, low validity cues will be relatively more under-utilized than high validity cues. Because high validity cues will, on average, give better predictions than low validity cues, their posterior probability of responsibility will increase, which in turn increases their weight in the final prediction. It is through this mechanism that the CBF explains the cue competition effects of Busemeyer et al. (1993a), who found that increasing the validity of one cue decreased the utilization of another cue. The CBF thus explains these cue competition effects as stemming from an erroneous assumption of cue competition in a non-competitive (linear) environment.

Global learning: Least Mean Squares network (LMS)

Despite the popularity of Bayesian models, it is not undisputed that learning involves updating probability distributions. Instead, learning may involve updating a single estimate of the validity weights based on prediction error. This learning process is implemented by the Least Mean Squares (LMS) model, which is a single-layer neural network with two input nodes (one for each cue) and a single output node for the criterion. The weights of the input-output connections reflect the inferred cue validity weights, and these are updated by means of the LMS or delta learning rule (e.g., Gluck & Bower, 1988). The LMS rule is structurally similar to the Kalman filter, but uses a learning rate parameter which is constant over time and does not depend on uncertainty. As such, the LMS rule can be taken as a simple approximation to the Kalman filter².

Local learning: Generalized Context Model (GCM)

The BLF, CBF and LMS network abstract information regarding the cue-criterion relations. Exemplar models, such as the context model (Medin & Schaffer, 1978) and Generalized Context Model (GCM, Nosofsky, 1986), offer an alternative which does not involve such abstraction. Exemplar models have been very successful in category learning, and more recently also in MCPL studies with a continuous criterion (e.g., Juslin et al., 2003). According to exemplar models, a

²It can be shown that, under certain conditions, the "learning rate" of the Kalman filter (the so-called Kalman gain, see Appendix A) converges to a constant (e.g. Durbin & Koopman, 2001, p.33). Under these conditions, the Kalman filter and LMS rule are equivalent in the limit. While convergence of the Kalman gain is not guaranteed for the DLM in general (in our case, the Kalman gain depends on the cue values, which vary from trial-to-trial and prevent convergence), the LMS rule can still be seen as an approximation to the Kalman filter.

response is made on the basis of a set of recollected exemplars (previously encountered cue patterns and corresponding criterion values). When asked to make a prediction, the cue pattern probes the recollection of stored exemplars. Each recollected exemplar is weighted according to the similarity of its cue values to the probe cue pattern. Finally, a response is made on the basis of a weighted average of recollected criterion values. In addition to similarity, we let the weight of exemplars depend on their recency. Memory decay of exemplars is necessary to track changes in the environment and we used two common functional forms of memory decay: an exponential (e.g. Nosofsky, Kruschke, & McKinley, 1992) and power function (e.g. Elliott & Anderson, 1995).

Local learning: Associative Learning Model (ALM)

Associative learning consists of gradually strengthening associations between cues and response alternatives. Busemeyer et al. (1997) and DeLosh et al. (1997) proposed a model of associative function learning, which they appropriately named the Associative Learning Model (ALM). The ALM was modelled after ALCOVE (Kruschke, 1992), a successful model of category learning. The ALM is implemented as an associative network with an input node for each unique cue pattern, and an output node for each unique criterion value. A probe cue pattern activates the input nodes according to a Gaussian generalization function, such that nodes receive more activation the more similar they are to the probe cue pattern. This spreaded activation of the input nodes is propagated to the output nodes according to a weighted additive rule, resulting in a spreaded activation of the output nodes. When making predictions, participants are assumed to choose one of the output nodes with a probability that reflects the relative activation of the output nodes. After observing the actual value of the criterion, the weights of the input-output connections are updated according to the same LMS rule used in the LMS model.

While the ALM and the LMS model are both implemented as single-layer artificial neural networks and apply the same learning rule, there are substantial differences. The LMS uses a single node for each cue and the criterion and is constrained to learn a linear relation between cues and criterion. The ALM uses many nodes for each cue and the criterion and therefore can also learn non-linear relations. Essentially, the ALM corresponds to a form of nonparametric density estimation (e.g., Silverman, 1986) of the conditional distribution of the criterion given cue patterns. As it makes no assumption about the shape of this distribution, it is a general model which will work in many situations. This generality comes at a price, however, as learning will be generally slower than in models which do make (approximately correct) assumptions about the distribution.

Results

We fitted the learning models to the data of Experiment 1 and 2. As for the DLMs, the learning models' free parameters (see Appendix C for details) were estimated under different equivalence assumptions. We fitted a version of each model in which the parameters were assumed identical between all participants, one in which the parameters were assumed to be identical within conditions, and one in which no between-participant equivalence was assumed. Full details of the models' fit are given in Appendix D.

Before comparing the different learning models, we will first discuss some basic findings for the different models. For all models, the results indicated the need to allow for individual variation in the model parameters, corroborating the evidence for individual variability in the DLM analyses. In addition, the global learning models showed that participants learned about the

Table 1: Best fitting learning models in Experiment 1 and 2

model	# par.	BIC	R^2	$n(\text{BIC})$	$n(R^2)$
Experiment 1					
BLF	112	75964	0.59	4	9
CBF	140	76791	0.55	4	8
LMS	84	76222	0.58	1	5
GCM	112	77413	0.53	1	3
ALM	140	74941	0.56	18	3
Experiment 2					
BLF	120	88736	0.36	4	8
CBF	150	88716	0.37	6	7
LMS	60	89268	0.33	0	2
GCM	120	89209	0.38	1	11
ALM	120	87680	0.35	19	2

Note: # par = number of freely estimated parameters (note that while all models had between 2 and 6 free parameters, as these were estimated for each individual participant, the total number of free parameters is relatively large), BIC = Bayesian Information Criterion, $n(\text{BIC})$ = number of participants for whom this was the best fitting model according to the BIC, $n(R^2)$ = number of participants for whom this was the best fitting model according to the R^2 , BLF = Bayesian Linear Filter, CBF = Competitive Bayesian Filter, LMS = Least Mean Squares network model, GCM = Generalized Context Model, ALM = Associative Learning Model (ALM).

cue validities at a different rate (the learning rate was identical only in the best fitting LMS model in Experiment 2). For the GCM, the results showed a clear superiority of the version with an exponential memory decay function (versions with exponential decay always fitted better than the corresponding versions with power decay). This is in contrast to the findings of Elliott and Anderson (1995), who found evidence for a power decay function. As their study involved a categorization task with gradually changing categories, their results may not have a direct bearing on the present task. Nonetheless, the difference is interesting, and merits further investigation. Note that the LMS model incorporates a form of exponential decay, as do the Bayesian filters (approximately). The best fitting version of the ALM in Experiment 1 was the one with a different generalization coefficient for the two cues. In Experiment 2, the best fitting version had an identical generalization coefficient for the two cues.

The fit measures of the best version of each model are collected in Table 1. Comparing the values of the BIC, we see that the ALM is the overall best fitting model in both experiments and thus provides the best description of the distribution of participants' responses. The superiority of the ALM may be partly due to the fact that it (correctly) takes participants' predictions to be integers between -100 and 100, rather than assuming predictions are continuous values not necessarily restricted to this range³. While the ALM had the best overall fit, it did not uniformly

³More technically, the ALM specifies a discrete probability distribution over the 201 possible responses, while the other models use continuous density functions over the whole real number line. For continuous density functions f , the likelihood l is usually defined as $l(R) = f(R)$, but to make comparable predictions to the ALM, we used the cumulative distribution functions $F(R) = \int_{-\infty}^R f(x)dx$, and computed the likelihood as $l(R) = F(R + .5) - F(R - .5)$.

fit best on an individual level. As Table 1 shows, the BLF and CBF fitted best for a number of participants, while there were a small number of participants for whom the GCM or LMS fitted best. When looking at the R^2 values, we see that the BLF outperformed the other models in Experiment 1. As the BLF has fewer free parameters than the ALM, its higher accuracy is not due to increased model complexity. In Experiment 2, the GCM performed best, using the same number of parameters as the ALM. Hence, we are left with a dilemma: the ALM appears to better capture the statistical properties of responses, while the BLF and GCM perform better in terms of their point predictions. We should note that, in contrast to the LMS and GCM, the reported R^2 values of the BLF, CBF and ALM are probably not as high as they could be, making the values for the BLF and CBF more impressive⁴.

The model fits are far from unequivocal. The relative fit of the ALM differs markedly depending on whether one looks at the BIC (best fitting model) or the R^2 (one of the worst fitting models). There is also a marked difference between Experiment 1 and 2: in the first experiment, the GCM was the worst fitting model, both in terms of the BIC and R^2 , while in the second experiment it outperformed all other models in terms of the R^2 , but only the LMS in terms of the BIC. Aggregating the data over the two experiments, we see that the ALM fits best according to the BIC while the BLF fits best according to the R^2 value⁵. With that in mind, we tentatively conclude that the BLF and ALM are the best models of participants' learning processes.

For a more qualitative comparison of the models' predictions, Figure 8 compares cue utilization as predicted by the models to the DLM estimates of participants' actual cue utilization. This shows that the differences between the models are mostly subtle. One difference is at the start of the task, where the CBF and GCM predict relatively quick learning, while the other models predict more gradual changes in cue utilization.

Experiment 3

Both the BLF and ALM fitted the data of Experiment 1 and 2 well, so it remains unclear whether learning is a local or global process. Experiment 3 was designed to resolve this ambiguity. As noted earlier, global learning models such as the BLE, CBF and LMS learn single functions which apply to the whole range of possible cue values. Local learning models, such as the ALM and GCM, effectively learn a large number of rules each applying to only a specific localized region in the space of possible cue values. The design of Experiment 3 was based on the observation that this distinction extends to the inference of change. For global models, a change detected in one region of the cue space will change the function globally, so that changes generalize to all possible cue values. For local models, a change in one region of the cue space

While this solves part of the problem by basing all computed likelihoods on discrete probability distributions, the continuous distributions assign positive mass to values outside the range of possible responses. This puts the ALM at an advantage in terms of likelihood values (which form the basis of the BIC).

⁴Due to the form of the response distributions of these models (Normal distributions with changing variance for the BLF and CBF, and a general discrete distribution for the ALM), maximum likelihood and least squares (maximum R^2) estimation are likely to give different parameter estimates. The response distributions of the LMS and GCM models are normal distributions with a constant variance, for which maximum likelihood and least squares estimation give identical results. Thus, for the LMS and GCM, maximum likelihood estimation results in minimum BIC and maximum R^2 values, while for the BLE, CBF and ALM, estimation only minimizes the BIC values. For this reason, maximum likelihood estimation will bias comparison on R^2 values towards the LMS and GCM

⁵The aggregated BIC values are 164860, 165708, 165671, 166782 and 162833, and the R^2 values are .45, .44, .43, .44 and .43, for the BLE, CBF, LMS, GCM and ALM respectively.

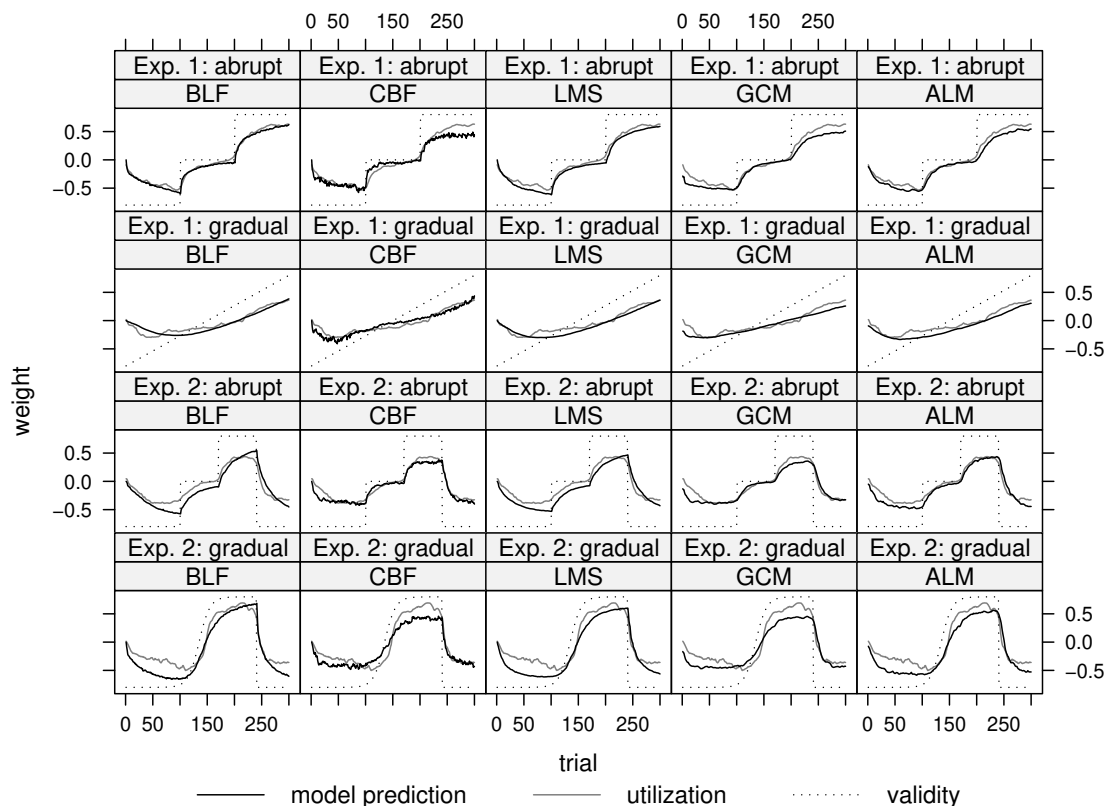


Figure 8. Average predicted utilization of cue 1 according to the BLF, CBF, LMS, GCM and ALM in Experiment 1 (top rows) and 2 (bottom rows). For the BLF and LMS, predicted cue utilization is identical to inferred cue validity. For the CBF, utilization is identical to the inferred cue utilization weighted by the (predicted) cue responsibility. For the GCM and ALM, we derived approximate predictions of cue utilization by fitting DLMs to the models’ predicted responses. Also shown are the average utilization curves as derived from the DLM, and the cue validity weights.

will not generalize to other regions. By following a change in cue validity with training on a restricted range of cue values, we should be able to discriminate between global and local forms of learning. While local learners should only learn to change their predictions to cue values in this restricted range, global learners should learn to also change their predictions to cue values outside this range. Testing participants on the whole range of cue values (without feedback) should thus provide a “model-free” test of global vs local learning.

In the previous two experiments, some participants seemed to have problems learning negative cue-criterion relations. This is consistent with results from previous research showing that negative relations are more difficult to learn than positive ones (e.g., Brehmer, 1973; Naylor & Clark, 1968). To increase the generalizability of the results, we included a Positive-Negative (PN) condition in which cue validity was initially positive before changing to a negative value, in addition to the Negative-Positive (NP) condition, where the cue validity changed from a negative

to a positive value.

Method

Participants. Forty students from University College London took part in the experiment. Participants were paid a basic fee (2 pounds) for their time, as well as an additional fee (between 0 and 5 pounds) based on their performance. Participants were randomly assigned to either the NP or PN condition. The two conditions did not differ in the distribution of sex (NP: 10 male, 10 female; PN: 8 male, 12 female; $\chi^2(1) = .10$, $p = .75$), or mean age (NP: $M = 20.00$, $SD = 0.92$; gradual $M = 19.60$, $SD = 0.88$; $t(38) = 1.41$, $p = .17$).

Materials. Apart from the inclusion of a test phase without outcome feedback and reducing the total number of trials to 220, the task was similar to that of the previous experiments. In the NP condition, the validity of cue 1 started at a negative value of $v_1 = -0.8$ for trials 1 to 100, and then changed to a positive value of $v_1 = 0.8$ for trials 101 to 220. In the PN condition, this pattern was reversed, starting with a positive validity of $v_{1t} = 0.8$ for trials 1 to 100 and then changing to a negative value of $v_1 = -0.8$ for trials 101 to 220. In both conditions, the validity of cue 2 was held constant at $v_2 = 0.4$. See Figure 5 for a graphical representation of the resulting cue validity patterns.

Precise details of the experimental design are given in Table 2. The task was divided into three phases: two learning phases and a test phase. Learning Phase 1 was similar to the task in the previous experiments. Participants were presented with cue values sampled from the range $[-100; 100]$, and asked to predict the (change in) share price of software company Mango. In Learning Phase 2 (trials 101 to 180), where the validity of cue 1 differed from Learning Phase 1, participants were presented with only a subset of the possible values of cue 1. For one group these values were in the ranges $[-50; 0]$ and $[50; 100]$, and for the other group in the ranges $[-100; -50]$ and $[0; 50]$. Thus, there were two ranges excluded from Learning Phase 2, an “interpolation” range, lying in between the two included ranges, and an “extrapolation” range, lying outside the two included ranges. In the Test Phase (trials 181 to 220), participants were asked to make predictions without outcome feedback. The cue validity was identical to that in the second learning phase.

In the Test Phase, participants were presented with critical and control cue patterns. The critical patterns had values for cue 1 which were outside the range used in Learning Phase 2, while the control patterns had values for cue 1 which were inside this range. Both critical and control patterns had been presented before at the end of Learning Phase 1 (and the control patterns also at the end of Learning Phase 2). Global learning models predict that participants give different predictions in the Test Phase compared to Learning Phase 2, and that this difference is comparable for the critical and control cue patterns. Local learning models predict that participants will make different predictions for the control cue patterns (as they have learned to do so during Learning Phase 2), but not, or to a much lesser extent, for the critical cue patterns.

Procedure. The procedure was similar to that of Experiment 1, but the instructions informed participants that they would also be asked to make predictions without outcome feedback. At the start of the test phase, participants were informed that they were now to make predictions without feedback, and that they could make their predictions as in the previous trials with feedback.

Table 2: Design of the task in Experiment 3. The values given are for half the participants; for the other half, critical and control patterns were interchanged, and the values of Cue 1 in Learning Phase 2 were randomly sampled from the ranges [-50;0] and [50;100].

<i>Learning Phase 1</i>	
Trial 1-80	Randomly generated cue patterns: <ul style="list-style-type: none"> • Cue 1 randomly sampled from the range [-100;100] • Cue 2 randomly sampled from the range [-100;100]
Trial 81-100	10 critical cue patterns: <ul style="list-style-type: none"> • Cue 1 = -45, -35, -25, -15, -5, 55, 65, 75, 85, 95 • Cue 2 randomly sampled from the range [-100;100] 10 control cue patterns: <ul style="list-style-type: none"> • Cue 1 = -95, -85, -75, -65, -55, 5, 15, 25, 35, 45 • Cue 2 randomly sampled from the range [-100;100]
<i>Learning Phase 2</i>	
Trial 101-160	Random generated cue patterns: <ul style="list-style-type: none"> • Cue 1 randomly sampled from the ranges [-100;-50] and [0;50] • Cue 2 randomly sampled from the range [-100;100]
Trial 161-180	10 random cue patterns (as above) 10 control cue patterns (repeated from Learning Phase 1)
<i>Test Phase</i>	
Trial 181-200	10 critical cue patterns (repeated from Learning Phase 1) 10 control cue patterns (repeated from Learning Phase 1 & 2)
Trial 201-220	10 new critical cue patterns (generated as in Learning Phase 1) 10 new control cue patterns (generated as in Learning Phase 1)

Results and discussion

The results are presented in largely the same order as earlier. We first discuss overall learning performance, followed by an investigation of the responses to the critical and control patterns. We then discuss cue utilization, followed by a comparison of the formal learning models. We end with the assessment of participants' insight. Due to a technical error, responses of nine participants on one critical trial (each with a different cue pattern) of the Test Phase were not recorded; these were treated as missing values.

Performance. The correlation between predictions and criterion values in the two learning phases ranged from -0.02 to 0.88 ($M = 0.45$, $SD = 0.24$) in the NP condition, and from 0.38 to 0.97 ($M = 0.64$, $SD = 0.19$) in the PN condition. For two participants in the NP condition (N10 and N13) the correlation was not significantly higher than 0, $t(178) = 1.36$, $p = .18$, and $t(178) = -0.31$, $p = 0.76$, respectively. The correlation was significantly higher than 0 for all other participants. As before, we report the results using data from all participants; unless stated otherwise, the same pattern of results was obtained after exclusion of N10 and N13.

A 2 (condition) \times 11 (block: 20 trials each) ANOVA with repeated measures on the second

factor and Greenhouse-Geisser correction was conducted for the MAD between predicted and actual values of the criterion⁶. This analysis showed a significant effect of block, $F(3.71, 140.91) = 6.90$, $MS_e = 330.42$, $p < .001$, which is due to a sharp increase in prediction error in the block where cue validity changed. There was also a significant interaction between condition and block, $F(3.71, 140.91) = 27.09$, $p < .001$, which is mainly due to smaller prediction error in the blocks where cue validity was positive ($M = 24.61$, $SD = 7.90$) compared to negative ($M = 41.06$, $SD = 3.99$). The main effect of condition was not significant (NP: $M = 34.80$, $SD = 32.71$; PN: $M = 30.43$, $SD = 32.69$), $F(1, 38) = 1.38$, $MS_e = 1519.13$, $p = 0.25$.

Critical and control trials. Figure 9 shows the average responses and optimal responses to the critical and control cue patterns. Each arrow represents the change in average prediction from Learning Phase 1 (arrow base) to the Test Phase (arrow tip). The results indicate that participants revised their responses after the change in cue validity, both for critical and control test trials. However, in general, participant's responses were more cautious (closer to 0) and changed less than optimal. In agreement with the analysis of the MAD, responses in Learning Phase 1 were more optimal in the PN condition than in the NP condition, while the reverse holds in the test phase. More importantly, there does not appear to be a consistent difference between the critical and control test trials in either condition. This was confirmed in a condition by trial type (control vs critical) by cue value ANOVA, with as dependent variable the ratio of the change in actual response to the change in optimal response from learning to test (i.e., the ratio of the length of the arrows for actual and optimal responses). This analysis only showed a significant effect of cue value, $F(19, 555) = 7.33$, $MS_e = 1.07$, $p < .001$, but no effects of trial type (all F 's < 1.50). The main effect of cue value was due to the "overshooting" for small values of cue 1 (values -5 and 5), as evident in Figure 9. After elimination of the responses to these two cue values, the main effect of cue value was no longer significant, $F(17, 555) = 1.14$, $MS_e = 0.40$, $p = .31$.

The absence of a trial type effect indicates that participants revised their responses to the critical test trials as much as their responses to control test trials. This is expected from global learning, where learned changes for the control cue patterns generalize directly to the critical test patterns. While local learning allows for some generalization of change to surrounding cue patterns, a key prediction is that this extrapolated change should be smaller in magnitude than learned change. As there was no indication of this effect, the results are more supportive of global than local learning.

Which cues were used? The pattern of results was similar to that obtained in the previous experiments (see Appendix B for details). The large majority ($n = 26$) of participants only relied on the two actual cues. Some participants ($n = 11$) also seemed to rely on the previous change in the criterion, and a small number of participants ($n = 4$) seemed to assume an interaction between the cues. Again, there was strong evidence for individual variation in the model parameters.

Cue utilization. The cue utilization profiles derived from the best fitting DLMs are given in Figure 10. As in the previous experiments, these figures show that participants' cue utilization adapted to the changing environment. As before there was a lot of individual variability; some participants (e.g., N7, N14, P4 and P13) adapted quickly and accurately to the changes, while

⁶For the test trials, the MAD refers to the absolute difference between predictions and the expected value of the criterion.

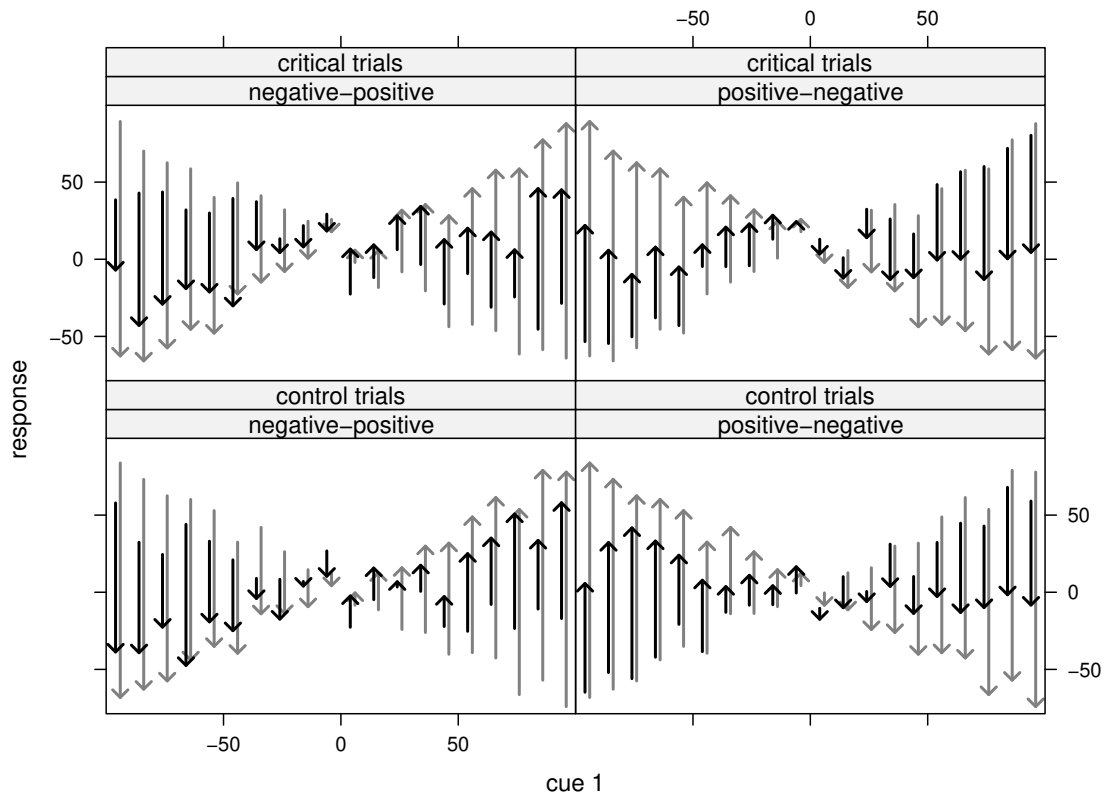


Figure 9. Mean actual and optimal responses to the critical and control test patterns. Arrows start at the response in Learning Phase 1 and end at the response in the Test Phase. Black arrows indicate the actual responses and grey arrows the optimal responses. Different test patterns are indicated by the value of the first cue, displaced slightly to separate actual and optimal responses. Note that these plots combine the results of two groups. In one group, the critical cue 1 values were between -50 and 0 and between 50 and 100, while in the other group, the critical cue 1 values were between -100 and -50 and between 0 and 50. This also holds for the control cue patterns.

others (e.g., N10 and P20) showed relatively little adaptation. Some participants (e.g., N8, N9 and N11) seem to have had initial difficulties in utilizing cue 1 according to its (negative) validity, but learned to utilize the cue in the direction of its positive validity in the second learning phase. In the PN condition, all participants showed positive utilization in Learning Phase 1. However, some participants (e.g., P3, P14 and P15) were apparently unable to later change this utilization in Learning Phase 2. Finally, some participants (e.g., P5 and P18) appeared to change their utilization from Learning Phase 2 to the Test Phase.

The average utilization is depicted in Figure 5. The conditions clearly differed in the utilization of cue 1, which was confirmed by a condition \times block (11 blocks of 20 trials each) ANOVA which showed a significant interaction between condition and block, $F(1.96, 74.45) = 65.31$, $MS_e = 0.41$, $p < .001$. The conditions also differed in the utilization of cue 2, as confirmed by a significant interaction between condition and block, $F(3.26, 124.05) = 9.13$, $MS_e = 0.08$, $p < .001$. In the NP condition, utilization of cue 2 slowly increased to its validity. In the PN condition, utilization deviated little from validity initially, but decreased in Learning Phase 2.

Learning models. As participants received no outcome feedback in the Test Phase, we assumed no learning would take place there, and therefore fitted the models to the data of the two learning phases, effectively using the data from the Test Phase as a cross-validation sample.

The complete results of the model fits are given in Appendix D. Table 3 contains the results for the best fitting version of each model (this selection was based all trials, i.e., the learning and test phases combined). For the learning phases we obtained similar results to Experiment 2: the BIC selected the ALM as the best fitting model, but the R^2 was highest for the GCM. Interestingly, although the GCM has the highest aggregated R^2 value, the BLF has the highest R^2 value for more individual participants than the GCM.

For present purposes, more important is how well the models predict responses in the Test Phase. As data from the test phase was not used in parameter estimation, the number of free parameters was set to 0 in calculating the BIC for this phase. As could be expected, the fit of all models is worse in the test phase than in the learning phases. The ALM even has a negative R^2 value, meaning that the model does a worse job in predicting participants' responses than the overall mean. The CBF outperforms the other models, both in terms of the BIC and R^2 . With regard to the distinction between global and local learning models, it is interesting to distinguish between the interpolation and extrapolation trials. Recall that these terms reflect the inter- and extrapolation of change to cue values encountered before the change, rather than the inter- and extrapolation of a stable cue-criterion relation to not yet encountered cue values. If learning is global and change extrapolated, we should expect the local learning models to perform especially poorly on the extrapolation trials, as the interpolation trials can benefit more from generalization from the surrounding cue values. This clearly holds for the ALM, which has the worst fit for the extrapolation trials, both in terms of the BIC and R^2 . However, the pattern is less clear for the GCM. While the GCM is outperformed on the extrapolation trials by all three global learning models in terms of the BIC, its R^2 value is only surpassed by the CBF. Moreover, in contrast to all other models, the R^2 values indicate better fit to the extrapolation trials than the interpolation trials. This is unexpected and it is unclear why the GCM would fit the extrapolation trials better than the interpolation trials.

To summarize, the results of the test phase indicate a superiority of the CBF, followed by the BLF, which performs better in the interpolation trials, but not as well in the extrapolation tri-

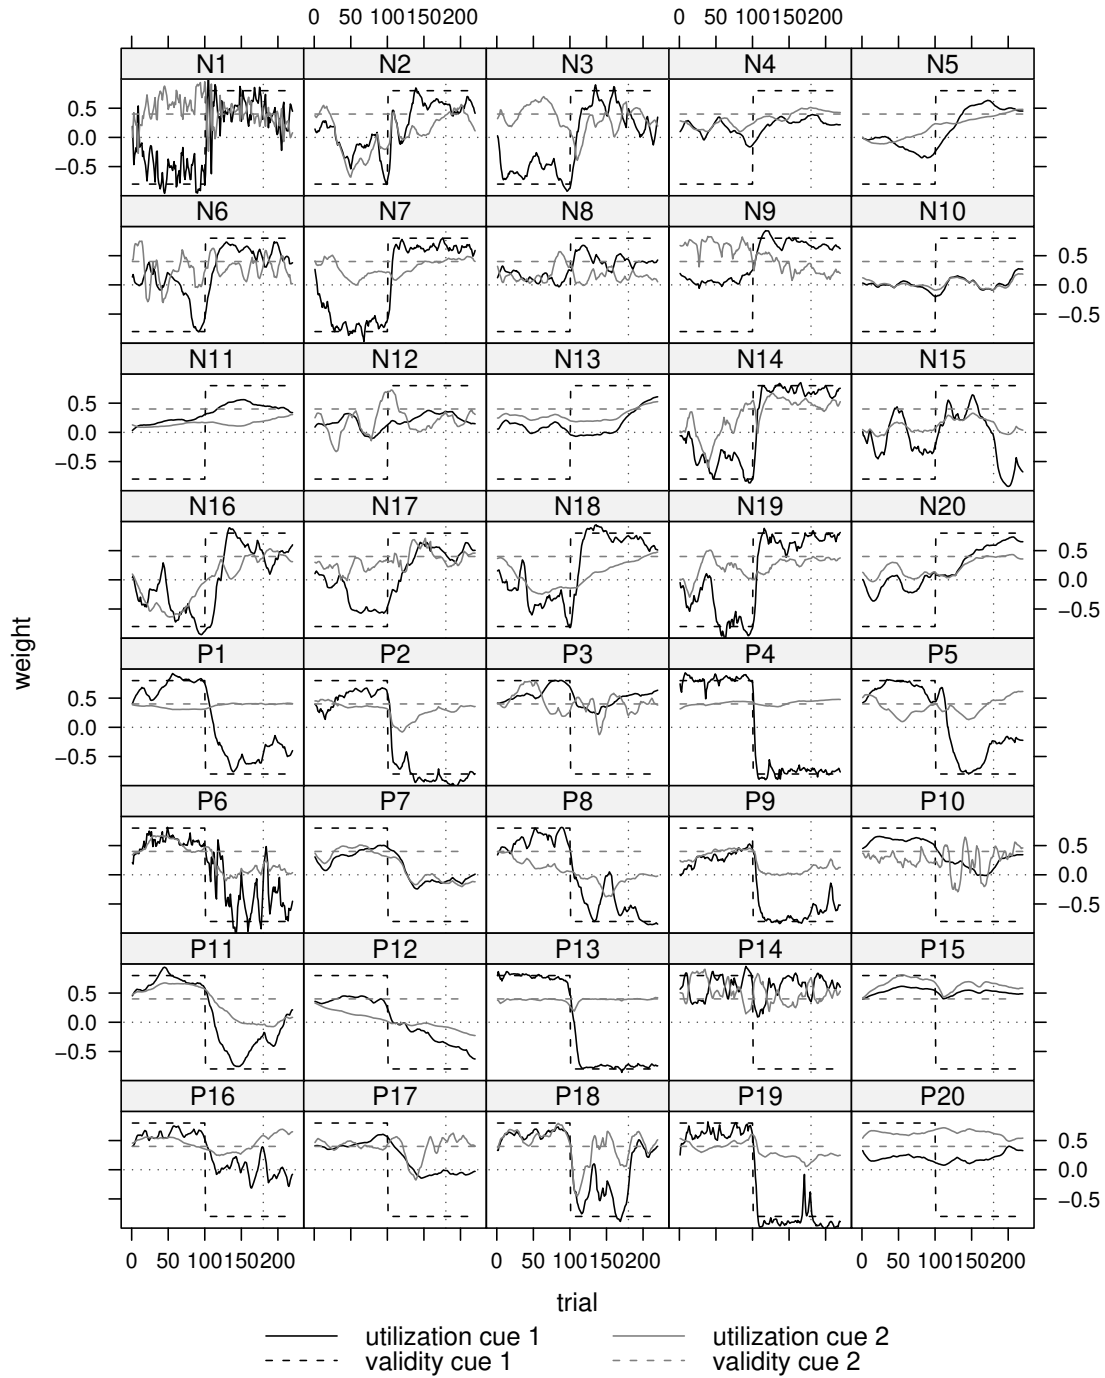


Figure 10. Cue utilization for participants in Experiment 3 in the Negative-Positive (N1-N20) and Positive-Negative (P1-P20) condition. Estimated utilization (solid lines) was computed with model C, apart from participants N17, P1, P13 (model I), P4, P5, P10, P12, P15, P18, N3, N9, N10 and N17 (model P) and N9 (model I+P). Validity (broken lines) represents actual cue validity. Horizontal dotted lines indicate zero cue utilization. Vertical dotted lines indicate the start of the test phase.

Table 3: Fit of the learning models in Experiment 3

model	# par.	BIC					R^2					$n(\text{BIC})$	$n(R^2)$
		learning	test			all	learning	test			all		
			inter.	extrap.	comb.			interp.	extrap.	comb.			
BLF	120	70711	3709	4163	7871	86383	0.42	0.41	0.16	0.24	0.41	7	13
CBF	200	71666	3732	4104	7836	87434	0.43	0.31	0.26	0.29	0.41	6	6
LMS	84	71416	3760	4172	7932	87264	0.42	0.29	0.05	0.13	0.38	0	2
GCM	160	71489	3840	4204	8044	87508	0.47	0.21	0.25	0.24	0.44	3	9
ALM	200	70090	3783	4322	8104	85940	0.44	0.29	-0.24	-0.06	0.39	24	10

Note: learning = Learning Phase 1 and 2, inter. = “interpolation” critical test trials, extrap. = “extrapolation” critical test trials, comb. = all critical test trials combined, all = all data from learning and test phases. See note in Table 1 for explanation of other abbreviations.

als. While the GCM performs quite well in the test phase in terms of the R^2 , it is outperformed by the CBF in this respect, and additionally by the BLF in terms of the BIC. Overall, participants’ inference of change appears more consistent with a global Bayesian learner than a local associative or exemplar learner.

Insight. The results of the insight questions are depicted in Figure 6. The results replicate those of the previous experiments. In general, the average ratings are consistent with the sign of the cue validity weights. A regression model with cue validity as single predictor showed a significant effect of cue validity on the ratings, $F(1, 239) = 102.3$, $p < .001$, $R^2 = 0.30$, but a similar model with cue utilization as predictor showed a stronger effect, $F(1, 239) = 123.4$, $p < .001$, $R^2 = 0.34$. And while cue utilization had an additional effect to cue validity, $F(1, 238) = 12.78$, $MS_e = 18.69$, $p < .001$, the reverse was not true, $F(1, 238) = 1.79$, $p = .18$.

The logistic regression model for the choice data with cue validity as predictor showed a significant effect of validity, $b = 1.55$ (S.E. = 0.48), $p < .01$, but did not fit the data very well, $\chi^2(119) = 155.39$, $p < .05$. The model with cue utilization as predictor showed a stronger effect of utilization, $b = 2.89$ (S.E. = 0.70), $p < .001$, and could not be rejected, $\chi^2(119) = 144.89$, $p = .053$. In addition, while adding cue utilization to the model with cue validity as predictor significantly improved model fit, $\chi^2(118) = 12.67$, $p < .001$, the reverse was not true, $\chi^2(118) = 2.18$, $p = .14$.

Thus, as in the previous experiments, participants showed good insight into what they learned, and we did not find evidence for implicit learning.

General discussion

We presented three experiments investigating how people learn to predict a criterion on the basis of two cues, where the relation between one of the cues and the criterion changed over time. The results showed that participants learned to predict the criterion reasonably well and in doing so were responsive to changes in the cue-criterion relations. Indeed, responses were quite subtly tuned to the dynamic changes in the environment, showing different patterns for gradual and abrupt changes (Experiments 1 and 2). Participants also showed good insight into the changing task structure. Comparing a number of formal learning models, we showed that participants’ behaviour was consistent with a Bayesian learner, who infers the parameters of a changing linear function by taking uncertainty into account, and an associative learner, who gradually strengthens the associations between cue patterns and response alternatives. The third experi-

ment showed that participants generalized a change learned for a particular range of cue values to the whole range of possible cue values. This is more consistent with global rule learning than local forms of learning.

Types of change

In the first two experiments, we compared learning in environments where changes in a cue-criterion relation were either gradual and smooth, or abrupt and discontinuous. The type of change appeared to have little effect on participants' overall prediction performance. In both types of environment, participants' cue utilization patterns reflected the changes in cue validity. Experiment 1 indicated that adaptation to gradual changes was slower than to abrupt changes. However, this result may have been due to problems in establishing an initial representation of the task. In the gradual change condition, cue validity changed from the outset, while validity was initially stable in the abrupt change condition. Keeping initial cue validity relatively stable in both conditions, we found no evidence that participants were slower to adapt to gradual changes in Experiment 2. This supports the hypothesis that an initial stable representation helps learning later changes. One potential caveat is that the gradual change in Experiment 2 was faster and therefore less different from the abrupt change than in Experiment 1. However, as a gradual curve fitted the cue utilization better than the abrupt curve in this condition, it seems unlikely that participants perceived the pattern of gradual changes as an abrupt one.

The second experiment also tested the hypothesis (e.g., Bouton, 2004; Elliott & Anderson, 1995; Redish et al., 2007) that abrupt changes will result in the formation of different task representations, rather than the adaptation of a single representation. We found no support for this hypothesis, as reinstatement of the environment's initial state did not result in an advantage for participants who previously encountered abrupt changes. Thus, abrupt changes did not appear to be learned in a qualitatively different way than gradual changes. In both cases, learning may entail the adaptation of a single representation to the changing environment. Of course, participants may come to expect changes of a similar kind as those encountered previously, so that a large abrupt change is more surprising after a period of relatively small, gradual changes than after a period of abrupt changes. If learning depends on surprise, an unexpected abrupt change may actually result in quicker adaptation (Yu & Dayan, 2005). This surprise would be on a higher level, stemming from a difference between expected and inferred change. Surprise is more commonly linked to changes in cue-criterion relations, rather than changes in the process underlying these changes. Detecting the latter type of change involves rather sophisticated inference. While the results are suggestive, how higher level change processes are reflected in the learning process requires further investigation. It is clear that these aspects can be learned (e.g., Behrens, Woolrich, Walton, & Rushworth, 2007). We can learn that some aspects of the environment remain stable, while others change, and that some changes are rapid, while others are slow. Addressing how we adapt the dynamics of our inferential processes to the environment is an important future challenge for cognitive science.

Implicit and explicit learning

Participants displayed good insight into general aspects of the cue-criterion relations. In particular, they were able to indicate the direction of the relation between each cue and criterion, as well as the relative strength of these relations. Moreover, they displayed awareness of the changes in cue validity. While participants' judgements were clearly related to cue validity,

the relation with cue utilization was even stronger. We take participants' cue utilization to reflect their inference of the task structure and conclude that, even though inferences may be prone to error (cue utilization and validity do not always match), participants were able to access and evaluate these inferences. This is in clear contrast to claims that MCPL tasks are learned implicitly.

Dynamic lens model analysis

We used a dynamic lens model (DLM) to obtain optimal estimates of participants' cue utilization on a trial-by-trial basis. The resulting fine-grained profiles of participants' learning process provided interesting results. First of all, many participants were clearly responsive to changes in the cue-criterion relations. In contrast to the results of earlier studies (e.g., Peterson et al., 1965), which indicated that participants were relatively slow in adjusting cue utilization to changes in validity, a large number of participants showed almost immediate changes in cue utilization following a change in validity. Secondly, the DLMs indicated large individual differences in cue utilization patterns. For instance, we could see that some participants appeared to have (initial) difficulties in utilizing a cue which was negatively related to the criterion. While previous MCPL studies have shown that negative relations are generally harder to learn than positive ones, we could also see that some participants seemed to suddenly overcome this difficulty, resulting in rapid changes in cue utilization. Such rapid changes may result from an attentional process, where the negative cue is first ignored, but suddenly attended to. Differences in attention may also underlie the cue competition effects displayed by some participants. While attentional processes were not explicitly incorporated into our analyses, there are different ways of doing so. One possibility is to allow the covariance matrix of utilization changes (Σ_u , see Appendix A) to change over time. Attention to a cue could be modelled by assigning that cue a relatively large variance, resulting in relatively rapid changes in utilization. A second possibility is to incorporate an "attentional gate" such that, on a given trial, participants can utilize a given cue or not. This trial specific utilization would be independent of the underlying latent utilization weights, such that participants can make sudden shifts between an effective utilization of 0, and one that is substantially different from 0. These options differ in important ways; the first is more directly related to the learning, the second more to the response process. While both can be relatively easily incorporated into the DLM, parameter estimation will become more involved. In the relatively simple case when participants only utilize a single cue at a time, we arrive at the Competitive Bayesian Filter. In general, we can no longer rely on the Kalman filter, and utilization estimates will no longer be analytically tractable. However, recent advances in approximate methods, such as Monte Carlo Markov Chain (e.g., Gilks, Richardson, & Spiegelhalter, 1996) and Particle Filters (e.g., Doucet, Freitas, & Gordon, 2001), mean that we can start exploring such extensions in future research.

Formal models of MCPL: global and local learning

To provide a more theoretically inspired account of participants' learning, we formulated five formal learning models. As noted in the introduction, model development for function learning and MCPL has lagged behind that for category learning. In addition, there has been little work on how these models function in dynamic environments. The present study aimed to fill this lacuna. We formulated two novel Bayesian MCPL models: the Bayesian Linear Filter

(BLF) and the Competitive Bayesian Filter (CBF), in addition to three previously proposed models, a Least Mean Squares (LMS) network model, the Generalized Context Model (GCM, Nosofsky, 1986), and the Associative Learning Model (ALM, Busemeyer et al., 1997).

The results of the model comparison were not conclusive. Apart from the LMS, all models performed best at some stage (the CBF outperforming the other models in the test phase of Experiment 3). For a final, overall evaluation, we aggregated the data from (the learning phases of) the three experiments. This shows that the ALM fitted best according to the BIC, while the R^2 measure indicates a superiority of the BLF⁷. The BLF's advantage is especially noteworthy since, in contrast to the LMS and GCM, the R^2 value was not directly maximised for this model and could potentially increase if parameters were estimated by a least squares procedure. As the BLF was also the second best performing model according to the BIC, it appears to be more of an "all-rounder" than the other models.

We were particularly interested in the difference between global and local learning. The results of the test phase in Experiment 3 seemed to support global over local learning models. The "model-free" test indicated that changes in one range of the cue values were generalized to other ranges. This extrapolation of change is a crucial prediction from global learning models. Local learning models can accommodate this extrapolation with a high level of generalization from experienced to new cue values. But this generalization should be weaker than directly learned changes. Yet we found no evidence that predictions were adjusted less for extrapolation cues. In addition, the global learning models showed superior fit to the test phase. While the first two experiments were ambivalent in their support for global and local learning models, the results of Experiment 3 clearly support global learning.

One caveat is that we tested the models in a linear task environment. Previous research indicates that while participants usually assume linearity, they can learn functions with more general shapes (e.g., Busemeyer et al., 1997; DeLosh et al., 1997; Kelley & Busemeyer, 2008). The BLF, like the other global learning models tested here, can only learn linear relations between cues and criterion and is unlikely to perform as well in highly nonlinear task environments. Unsurprisingly, models without linearity constraints, such as the GCM and ALM, perform better in those environments (Juslin et al., 2003; McDaniel & Busemeyer, 2005). Juslin et al. (2003, 2008) argued that people switch from rule-based to exemplar-based learning to cope with such environments. While there is no need to invoke multiple learning systems, the results do indicate that linear models such as the BLF may need to be extended to capture learning in more general MCPL tasks. While previous studies focussed on function learning (learning the relation between a single cue and the criterion), a recent study by Kelley and Busemeyer (2008) reports evidence that people can learn non-linear and non-additive functions of two cues.

Another potential criticism is that our conclusions are based on a limited number of learning models. Although compared to category learning, the literature contains relatively few models of MCPL, it would have been infeasible to compare all of them. We restricted our attention to models with relatively dissimilar characteristics. Although the BLF and LMS model are relatively similar, we included the latter to assess the evidence for a Bayesian learning mechanism which effectively adjusts the learning rate according to uncertainty. Interestingly, the results do support the existence of such a mechanism of learning rate adjustment. In addition, we tested specific versions of each model and it is possible that other variants would have performed better. For

⁷The aggregated BIC values were 251515, 253479, 253098, 254559 and 249110, and the R^2 values 0.441, 0.430, 0.415, 0.438, and 0.418, for the BLF, CBF, LMS, GCM and ALM respectively.

instance, we did not fully estimate participants' prior expectations regarding cue validity in the global learning models (i.e., we set the mean of the prior cue validity distribution to 0 in the BLF and CBF, and the starting weights to 0 in the LMS). Also, we tested a relatively simple version of the ALM which did not include the additional linear inter- and extrapolation component of the EXtrapolation Association Model (EXAM, DeLosh et al., 1997). This component allows the model to better capture participants' extrapolation of a function to unknown cue values. As our experiments did not involve an explicit extrapolation test, we deemed it unnecessary here; while the critical test trials in Experiment 3 required participants to generalize a change in the function to certain cue values, crucially, they had encountered these cue values before the change and did not extrapolate to unknown cue values. The tested version of the ALM also did not include the ability to learn dimensional attention weights; as in the GCM, we assumed that the generalization gradient for each cue remained the same throughout the task. While not incorporated into the original version of the ALM or EXAM, learning these generalization gradients (or dimensional attention weights) is an important aspect of ALCOVE (Kruschke, 1992), the category learning model on which the ALM was based. Whether it improves the performance of the ALM in multiple cue tasks will require further investigation. Similarly, future work can focus on extensions of the other models. For example, while the current versions of the BLF and CBF assume changes in cue validity have a constant (co-)variance throughout the task, incorporating mechanisms to infer changes in the magnitude of cue validity changes may improve upon these models.

Cue competition

Research on MCPL, particularly that conducted in the Brunswikian paradigm, typically assumes that cues affect predictions independently. Not only does this assumption underlie lens model analysis, there is evidence that participants have difficulties in learning to perform MCPL tasks when cues interact (e.g., Busemeyer et al., 1997), which may indicate that people naturally assume independent and additive cue effects. There are however findings suggesting that rather than being independent, cues have competing effects on predictions (Birnbaum, 1976; Busemeyer et al., 1993a). The strongest evidence for this cue competition was found by Busemeyer et al. (1993a), and it was noted (Busemeyer, Myung, & McDaniel, 1993b) that cue competition poses a serious problem for a large class of learning models. Our analysis of individual cue utilization patterns also showed signs of cue competition effects, albeit for a minority of participants. To further investigate evidence for cue competition, we formulated a novel model, the Competitive Bayesian Filter, which explicitly predicts cue competition effects. According to this model, cue competition results from an incorrect assumption regarding the structure of the environment, according to which only a single cue at a time can determine the criterion. The CBF learns optimally under this (invalid) assumption. The CBF fitted the data reasonably well, and it offered the best description of the behaviour of a number of participants. Although cue competition thus characterises a proportion of learners, our lens model analysis clearly indicates it is not a universal characteristic of human multiple cue learning.

Individual variation

One striking aspect of our results is the large variation between participants. Even in a relatively simple two-cue task, the analyses of participants' cue utilization show a wide range in learning dynamics. The usual practice of analysing aggregate data, or averaging over individual

analyses, can hide the richness in idiosyncratic behaviour. Some of the individual variation may reflect differences in sensitivity to change in the task structure. In our experiments, as in many real-life situations, participants were not explicitly alerted to the possibility of changes in cue-criterion relations. Different results could have been obtained if they were informed beforehand of the possibility of change. But our analyses show that most participants did learn that the environment changed. Moreover, additional research not reported here, in which participants were informed of the possible changes, showed similar results to those obtained here.

The idiosyncratic cue utilization patterns could alternatively reflect qualitative differences in learning processes. In other words, do participants learn in qualitatively different ways, or are their differences a matter of degree? The finding that the best fitting model varies over individuals seems to support the former option. However, as illustrated in Figure 8, differences in the models' predictions are often subtle, even when considering the full course of their learning rather than asymptotic performance. Small differences do occur, such as the relatively quick adaptation of the GCM and CBF in comparison to the other models. By incorporating changes in the environment, model discriminability may be enhanced, especially when controlling the cue patterns as in Experiment 3.

The difficulty in model discrimination necessitates the use of highly sensitive tests. For instance, models may differ mostly in certain periods, such as the start of the task (where the CBF and GCM adapt quicker than the other models, see Figure 8) or after a change in task structure. It will then be useful to focus experiments mainly there. And if models differ mostly in their predictions for certain cue patterns, then these patterns should be emphasised in the task. One complication in the design of such discriminative tests is that they are often highly sensitive to changes in model parameters. We found evidence of significant individual variation in parameter values, and therefore a sensitive test for one participant may not be a sensitive test for another. We are currently working on methods to resolve this problem, providing adaptive tests for model discrimination.

References

- Anderson, N. H. (1981). *Foundations of information integration theory*. New York: Academic Press.
- Baddeley, R. J., Ingram, H. A., & Miall, R. C. (2003). System identification applied to a visuomotor task: Near-optimal human performance in a noisy changing task. *Journal of Neuroscience*, *23*, 3066–3075.
- Behrens, T. E. J., Woolrich, M. W., Walton, M. E., & Rushworth, M. F. S. (2007). Learning the value of information in an uncertain world. *Nature Neuroscience*, *10*, 1214–1221.
- Berry, D. C., & Broadbent, D. E. (1984). On the relationship between task performance and associated verbalizable knowledge. *Quarterly Journal of Experimental Psychology*, *36*, 209–231.
- Birnbaum, M. H. (1976). Intuitive numerical prediction. *American Journal of Psychology*, *89*, 417–429.
- Bouton, M. E. (2004). Context and behavioral processes in extinction. *Learning and Memory*, *11*, 485–494.
- Brehmer, B. (1973). Single-cue probability learning as a function of the sign and magnitude of the correlation between cue and criterion. *Organizational Behavior and Human Performance*, *9*, 377–395.
- Brehmer, B. (1994). The psychology of linear judgment models. *Acta Psychologica*, *87*, 137–154.
- Brunswik, E. (1955). Representative design and probabilistic theory in a functional psychology. *Psychological Review*, *62*(3), 193–217.
- Busemeyer, J. R., Byun, E., DeLosh, E. L., & McDaniel, M. A. (1997). Learning functional relations based on experience with input-output pairs by humans and artificial neural networks. In K. Lamberts & D. Shanks (Eds.), *Knowledge, concepts and categories* (pp. 405–437). Psychology Press.

- Busemeyer, J. R., Myung, I. J., & McDaniel, M. A. (1993a). Cue competition effects: Empirical tests of adaptive network learning models. *Psychological Science*, 4, 190–195.
- Busemeyer, J. R., Myung, I. J., & McDaniel, M. A. (1993b). Cue competition effects: Theoretical implications for adaptive network learning models. *Psychological Science*, 4, 196–202.
- Cheng, S., & Sabes, P. N. (2007). Calibration of visually guided reaching is driven by error-corrective learning and internal dynamics. *Journal of Neurophysiology*, 97, 3057–3069.
- Chhabra, M., & Jacobs, R. A. (2006). Near-optimal human adaptive control across different noise environments. *Journal of Neuroscience*, 26, 10883–10887.
- Cooksey, R. W. (1996). The methodology of social judgement theory. *Thinking and Reasoning*, 2, 141–173.
- Dayan, P., Kakade, S., & Montague, P. R. (2000). Learning and selective attention. *Nature Neuroscience*, 3, 1218–1223.
- Dayan, P., & Long, T. (1998). Statistical models of conditioning. In M. I. Jordan, M. J. Kearns, & S. A. Solla (Eds.), *Advances in neural information processing systems 10: Proceedings of the 1997 conference*. Cambridge, MA: MIT Press.
- DeLosh, E. L., Busemeyer, J. R., & McDaniel, M. A. (1997). Extrapolation: the sine qua non for abstraction in function learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 23, 968–986.
- Dempster, A. P., Laird, N. M., & Rubin, D. B. (1977). Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society, Series B*, 39, 1–38.
- DeShon, R. P., & Alexander, R. A. (1996). Goal setting effects on implicit and explicit learning of complex tasks. *Organizational Behavior and Human Decision Processes*, 65, 18–36.
- Doucet, A., Freitas, N. de, & Gordon, N. (Eds.). (2001). *Sequential monte carlo methods in practice*. New York: Springer.
- Dudycha, A. L., Dumoff, M. G., & Dudycha, L. W. (1973). Choice behavior in dynamic environments. *Organizational Behavior and Human Performance*, 9, 328–338.
- Durbin, J., & Koopman, S. J. (2001). *Time series analysis by state space methods*. Oxford: Oxford University Press.
- Elliott, S. W., & Anderson, J. R. (1995). Effect of memory decay on predictions from changing categories. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 21, 815–836.
- Evans, J. S. B. T., Clibbens, J., Cattani, A., Harris, A., & Dennis, I. (2003). Explicit and implicit processes in multicue judgment. *Memory & Cognition*, 31, 608–618.
- Ghahramani, Z., & Hinton, G. E. (2000). Variational learning for switching state-space models. *Neural Computation*, 12(4), 831–864.
- Gilks, W. R., Richardson, S., & Spiegelhalter, D. J. (Eds.). (1996). *Markov chain Monte Carlo in practice*. Boca Raton: Chapman and Hall.
- 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. A., Shohamy, D., & Myers, C. (2002). How do people solve the “weather prediction task”? Individual variability in strategies for probabilistic category learning. *Learning & Memory*, 9, 408–418.
- Grant, D. A., & Berg, E. (1948). A behavioral analysis of degree of reinforcement and ease of shifting to new responses in a Weigl-type card-sorting problem. *Journal of Experimental Psychology*, 38, 404–411.
- Griffiths, T. L., & Tenenbaum, J. B. (2005). Structure and strength in causal induction. *Cognitive Psychology*, 51, 354–384.
- Hammond, K. R., Wilkins, M. M., & Todd, F. J. (1966). A research paradigm for the study of interpersonal learning. *Psychological Bulletin*, 65, 221–232.
- Harries, C., Evans, J. S. B. T., & Dennis, I. (2000). Measuring doctors’ self-insight into their treatment decisions. *Applied Cognitive Psychology*, 14, 455–477.
- Harvey, A. C. (1989). *Forecasting, structural time series models and the Kalman filter*. Cambridge: Cambridge University Press.

- Juslin, P., Karlsson, L., & Olsson, H. (2008). Information integration in multiple cue judgment: A division of labor hypothesis. *Cognition*, *106*, 259-298.
- Juslin, P., Olsson, H., & Olsson, A.-C. (2003). Exemplar effects in categorization and multiple-cue judgment. *Journal of Experimental Psychology: General*, *132*, 133-156.
- Kakade, S., & Dayan, P. (2002). Acquisition and extinction in autoshaping. *Psychological Review*, *109*, 533-544.
- Kalish, M. L., Lewandowsky, S., & Kruschke, J. K. (2004). Population of linear experts: knowledge partitioning and function learning. *Psychological Review*, *111*, 1072-1099.
- Kalman, R. E. (1960). A new approach to linear filtering and prediction problems. *Transactions of the American Society of Mechanical Engineers, Series D, Journal of Basic Engineering*, *82*, 35-45.
- Kalman, R. E., & Bucy, R. S. (1961). New results in linear filtering and prediction theory. *Transactions of the American Society of Mechanical Engineers, Series D, Journal of Basic Engineering*, *83*, 95-108.
- Kelley, H., & Bussemeyer, J. (2008). A comparison of models for learning how to dynamically integrate multiple cues in order to forecast continuous criteria. *Journal of Mathematical Psychology*, *52*, 218-240.
- Kelley, H., & Friedman, D. (2002). Learning to forecast price. *Economic Inquiry*, *40*, 556-573.
- Knowlton, B. J., Mangels, J. A., & Squire, L. R. (1996). A neostriatal habit learning system in humans. *Science*, *273*, 1399-1402.
- Körding, K. P., & Wolpert, D. M. (2004). Bayesian integration in sensorimotor learning. *Nature*, *427*, 244-247.
- Kruschke, J. K. (1992). ALCOVE: an exemplar-based connectionist model of category learning. *Psychological Review*, *99*, 22-44.
- Lagnado, D. A., Newell, B. R., Kahan, S., & Shanks, D. R. (2006). Insight and strategy in multiple cue learning. *Journal of Experimental Psychology: General*, *135*, 162-183.
- Lindberg, L.-A., & Brehmer, B. (1976). Transfer in single-cue probability learning. *Organizational Behavior and Human Performance*, *16*, 177-192.
- Lovibond, P. E., & 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.
- Mather, G. (2006). *Foundations of perception*. Psychology Press.
- McDaniel, M. A., & Bussemeyer, J. R. (2005). The conceptual basis of function learning and extrapolation: comparison of rule-based and associative-based models. *Psychonomic Bulletin & Review*, *12*, 24-42.
- Medin, D., & Schaffer, M. (1978). Context theory of classification learning. *Psychological Review*, *85*, 207-238.
- Milner, B. (1963). Effects of brain lesions on card sorting. *Archives of Neurology*, *9*, 90-100.
- Murphy, K. P. (1998). *Switching Kalman filters* (Tech. Rep. No. 98-10). Compaq Cambridge Research Lab.
- Naylor, J. C., & Clark, R. D. (1968). Intuitive inference strategies in interval learning as a function of magnitude and sign. *Organizational Behavior and Human Performance*, *3*, 378-399.
- Newell, B. R., Lagnado, D. A., & Shanks, D. R. (2007). Challenging the role of implicit processes in probabilistic category learning. *Psychonomic Bulletin & Review*, *14*, 505-511.
- Nisbett, R. E., & Wilson, T. D. (1977). Telling more than we can know: Verbal reports on mental processes. *Psychological Review*, *84*, 231-259.
- Nosofsky, R. M. (1986). Attention, similarity, and the identification-categorization relationship. *Journal of Experimental Psychology: General*, *115*, 39-61.
- Nosofsky, R. M., Kruschke, J. K., & McKinley, S. C. (1992). Combining exemplar-based category representations and connectionist learning rules. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *18*, 211-233.
- Oaksford, M., & Chater, N. (2007). *Bayesian rationality*. Oxford: Oxford University Press.
- Peterson, C. R., Hammond, K. R., & Summers, D. A. (1965). Multiple probability-learning with shifting weights of cues. *American Journal of Psychology*, *4*, 660-663.

- R Development Core Team. (2006). R: A language and environment for statistical computing [Computer software manual]. Vienna, Austria. Available from <http://www.R-project.org> (ISBN 3-900051-07-0)
- Redish, A. D., Jensen, S., Johnson, A., & Kurth-Nelson, Z. (2007). Reconciling reinforcement learning models with behavioral extinction and renewal: implications for addiction, relapse, and problem gambling. *Psychological Review*, *114*, 784–805.
- Rescorla, R. A. (1999). Partial reinforcement reduces the associative change produced by nonreinforcement. *Journal of Experimental Psychology: Animal Behavior Processes*, *25*, 403–414.
- Ruffner, J. W., & Muchinsky, P. M. (1978). The influence of shifting cue validity distributions and group discussion feedback on multiple cue probability learning. *Organizational Behavior and Human Performance*, *21*, 189–208.
- Schwarz, G. (1978). Estimating the dimension of a model. *Annals of Statistics*, *6*, 461–464.
- Seger, C. A. (1994). Implicit learning. *Psychological Bulletin*, *115*, 163–196.
- Shumway, R. H., & Stoffer, S. (1982). An approach to time series smoothing and forecasting using the EM algorithm. *Journal of Time Series Analysis*, *3*, 253–264.
- Silverman, B. W. (1986). *Density estimation for statistics and data analysis*. London: Chapman and Hall.
- Slovic, P., & Lichtensteln, S. (1971). Comparison of bayesian and regression approaches to the study of information processing in judgment. *Organizational Behavior and Human Performance*, *6*, 649–744.
- Speekenbrink, M., Channon, S., & Shanks, D. R. (2008). Learning strategies in amnesia. *Neuroscience and Biobehavioral Reviews*, *32*, 292–310.
- Speekenbrink, M., & Shanks, D. R. (2008). Through the looking glass: A dynamic lens model approach to learning in MCPL tasks. In N. Chater & M. Oaksford (Eds.), *The probabilistic mind: Prospects for rational models of cognition* (pp. 409–429). Oxford: Oxford University Press.
- Summers, D. A. (1969). Adaptation to change in multiple probability tasks. *American Journal of Psychology*, *8*, 235–240.
- Sutton, R. S. (1992). Gain adaptation beats least squares? In *Proceedings of the 7th Yale workshop on adaptive and learning systems* (pp. 161–166).
- Tenenbaum, J. B., Griffiths, T. L., & Kemp, C. (2006). Theory-based bayesian models of inductive learning and reasoning. *Trends in Cognitive Science*, *10*, 309–318.
- West, M., & Harrison, J. (1997). *Bayesian forecasting and dynamic models* (2nd ed.). New York: Springer.
- Wigton, R. S. (1996). Social judgment theory and medical judgment. *Thinking and Reasoning*, *2*, 175–190.
- Wolpert, D. M., Ghahramani, Z., & Jordan, M. I. (1995). An internal model for sensorimotor integration. *Science*, *269*, 1880–1882.
- Wu, L. S.-Y., Pai, J. S., & Hosking, J. R. M. (1996). An algorithm for estimating parameters of state-space models. *Statistics and Probability Letters*, *28*, 99–106.
- 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.
- Yu, A. J., & Dayan, P. (2005). Uncertainty, neuromodulation, and attention. *Neuron*, *46*, 681–692.

Appendix A

Cue utilization estimation by the Kalman filter/smoothing

In this appendix, we give details of the Dynamic Lens Model (DLM) and the Kalman filter and smoother techniques for its estimation. We also give a brief description of the Expectation-Maximisation (EM) algorithm to estimate the model parameters. More general and extensive treatments can be found in Durbin and Koopman (2001), Harvey (1989), or West and Harrison (1997). The presentation is geared towards the current situation, and the equations presented are not optimal from a computational viewpoint. For the actual implementation, we used the

more numerically stable computations described in Wu, Pai, and Hosking (1996).

Dynamic Lens Model (DLM)

In the DLM, we model a response on trial t , denoted as R_t , as a function of the cue values and utilization weights as follows:

$$R_t = \mathbf{x}_t \mathbf{u}_t + e_t, \quad e_t \sim N(0, \sigma_r), \quad (2)$$

where $\mathbf{x}_t = (x_{1t}, \dots, x_{mt})$ is the $1 \times m$ vector containing the cue values, and $\mathbf{u}_t = (u_{1t}, \dots, u_{mt})^\top$ the $m \times 1$ vector containing the utilization weights (\top denotes the matrix transpose). Note that this is just a multiple regression model written in matrix form; the response is the dependent variable, the cues the independent variables, and the utilization weights the regression coefficients. The right hand part of Equation 2 states that the error term, e_t , is a Normally distributed variable with zero mean and a standard deviation σ_r , which reflects how variable responses are around the model predictions. We can use a similar model for the task environment, replacing the response by the criterion value Y_t and the cue utilization weights by the cue validity weights \mathbf{v}_t . The cue utilization and/or validity weights are allowed to vary over trials/time and these changes are assumed to follow a first-order Markov process, so that the utilization (validity) at a particular time point only depends on the utilization (validity) at the previous time point (see also Figure 2). The distribution of the cue utilization (validity) at time $t + 1$, conditional upon the utilization (validity) at time t , is specified by a transition model. For the cue utilization weights, we use the following transition model

$$\mathbf{u}_{t+1} = \mathbf{u}_t + \mathbf{d}_t, \quad \mathbf{d}_t \sim N(0, \Sigma_u), \quad (3)$$

i.e., the utilization at time $t + 1$ has a (multivariate) Normal distribution centred on the utilization at time t . This transition model, in which the probability of an increase is identical to the probability of a corresponding decrease, is also known as a random walk. The covariance matrix Σ_u effectively determines the magnitude of the cue utilization changes: if the diagonal values of this matrix (the variances) are relatively large, utilization will undergo relatively large changes from one trial to the next, while small values result in relatively stable utilization over trials. The covariances in Σ_u reflect the inter-dependency between changes in the utilization of the cues. The specification of the DLM is completed by a prior distribution of utilization weights, for which we use a multivariate Normal distribution:

$$\mathbf{u}_0 \sim N(\boldsymbol{\mu}_0, \Sigma_0). \quad (4)$$

The mean vector ($\boldsymbol{\mu}_0$) and covariance matrix (Σ_0) of the prior distribution reflect participants' initial guess regarding the effect of the cues. The higher the variances in Σ_0 , the more uncertain this initial guess is, and the quicker participants will adapt their utilization in the initial trials. The prior distribution thus determines both the starting point of utilization and the magnitude of initial changes in utilization. Its effect on later utilization decreases, and eventually it "washes out" completely.

Given values for the model parameters, $\boldsymbol{\mu}_0$, Σ_0 , Σ_u and σ_r , the remaining task is to estimate the values of \mathbf{u}_t for $t = 1, \dots, T$. For this, we use the Kalman filter and smoother. The Kalman filter computes an estimate of \mathbf{u}_t based on the responses thus far (i.e., R_1, R_2, \dots, R_t). It is an on-line estimation technique, useful when observations are made sequentially and we

need an estimate at each time point. As later observations provide useful information about previous values of \mathbf{u}_t , they should be used if available. The Kalman smoother does just that, and provides estimates of \mathbf{u}_t taking all observations into account (i.e., R_1, R_2, \dots, R_T , where T denotes the total number of trials). As we estimated cue utilization after completion of all trials, we relied the Kalman smoother. Effectively, the Kalman smoother updates estimates from the Kalman filter on the basis of later observations. Below, we provide details of both the Kalman filter and smoother techniques. The estimation routines were implemented in custom software written in R, and an R package is available from the first author's website which can be accessed via http://www.ucl.ac.uk/psychlangsci/staff/cpb-staff/m_speekenbrink.

Filtering

Starting from an initial distribution $\mathbf{u}_0 \sim N(\boldsymbol{\mu}_0, \boldsymbol{\Sigma}_0)$, *filtering* consists of the recursive estimation of $\mathbf{u}_{t=k}$ based on the observations thus far (i.e., R_t and \mathbf{x}_t , $t = 1, \dots, k$). In the following, let $R_{j:k} = (R_j, R_{j+1}, \dots, R_k)$ and similarly $\mathbf{x}_{j:k} = (\mathbf{x}_j, \mathbf{x}_{j+1}, \dots, \mathbf{x}_k)$. From a Bayesian perspective, filtering consists of computing the posterior distribution $p(\mathbf{u}_t | R_{1:t}, \mathbf{x}_{1:t})$. The cue vectors \mathbf{x}_t are assumed fixed and we will not explicitly conditionalize on \mathbf{x}_t in the remainder.

For the DLM specified above, all relevant distributions are (multivariate) Normal. From multivariate Normal theory, it can be shown (e.g., Durbin & Koopman, 2001; West & Harrison, 1997) that the *filtering* distribution is

$$p(\mathbf{u}_t | R_{1:t}) = N(\mathbf{m}_t, \mathbf{C}_t), \quad (5)$$

the *state prediction* distribution is

$$p(\mathbf{u}_{t+1} | R_{1:t}) = N(\mathbf{m}_t, \mathbf{P}_t), \quad (6)$$

and the *forecast* distribution is

$$p(R_{t+1} | R_{1:t}) = N(\mathbf{x}_{t+1} \mathbf{m}_t, V_t). \quad (7)$$

The Kalman filter is an effective algorithm to iteratively compute the parameters (means and (co)variances) of these distributions. The mean \mathbf{m}_t of the filtering distribution can be computed as

$$\mathbf{m}_t = \mathbf{m}_{t-1} + \mathbf{K}_t (R_t - \mathbf{x}_t \mathbf{m}_{t-1}), \quad (8)$$

which is initialized at $\mathbf{m}_0 = \boldsymbol{\mu}_0$. The covariance of the filtering distribution can be computed as

$$\mathbf{C}_t = \mathbf{P}_t - \mathbf{K}_t V_t \mathbf{K}_t^\top, \quad (9)$$

and is initialized at $\mathbf{C}_0 = \boldsymbol{\Sigma}_0$. The *Kalman gain* term is

$$\mathbf{K}_t = \mathbf{P}_t \mathbf{x}_t^\top / V_t. \quad (10)$$

The covariance matrix of the state prediction distribution is computed as

$$\mathbf{P}_t = \mathbf{C}_{t-1} + \boldsymbol{\Sigma}_u. \quad (11)$$

Finally, the variance V_t of the forecast distribution is

$$V_t = \mathbf{x}_t \mathbf{P}_t \mathbf{x}_t^\top + \sigma_r^2. \quad (12)$$

Smoothing

The Kalman filter provides optimal estimates of \mathbf{u}_t at each time t , using only the observations made up to that time. As mentioned previously, at a later time $t + k$, there is more information available to estimate \mathbf{u}_t , so the filtering estimates are (not necessarily) optimal in light of this new evidence. Smoothing consists of estimating \mathbf{u}_t on the basis of *all* observations $R_{1:T}$.

The smoothing distribution is

$$p(\mathbf{u}_t | R_{1:T}) = N(\mathbf{m}'_t, \mathbf{C}'_t). \quad (13)$$

The mean \mathbf{m}'_t of the smoothing distribution can be computed as

$$\mathbf{m}'_t = \mathbf{m}_t + \mathbf{P}_t \mathbf{q}_t, \quad (14)$$

where

$$\mathbf{q}_{t-1} = \mathbf{x}_t^\top V_t^{-1} (R_t - \mathbf{x}_t \mathbf{m}_t) + (\mathbf{I} - \mathbf{K}_t \mathbf{x}_t)^\top \mathbf{q}_t, \quad (15)$$

initialized at $\mathbf{q}_T = 0$, and \mathbf{I} is an identity matrix. The covariance matrix \mathbf{C}'_t of the smoothing distribution can be computed as

$$\mathbf{C}'_t = \mathbf{P}_t - \mathbf{P}_t \mathbf{M}_{t-1} \mathbf{P}_t, \quad (16)$$

where

$$\mathbf{M}_{t-1} = \mathbf{x}_t^\top V_t^{-1} \mathbf{x}_t + (\mathbf{I} - \mathbf{K}_t \mathbf{x}_t)^\top \mathbf{M}_t (\mathbf{I} - \mathbf{K}_t \mathbf{x}_t), \quad (17)$$

which is initialized at $\mathbf{M}_T = 0$.

Parameter estimation

The Kalman filter and smoother assume that parameters $\boldsymbol{\mu}_0$, $\boldsymbol{\Sigma}_0$, $\boldsymbol{\Sigma}_u$ and σ_r are known. When this is not the case, maximum likelihood estimates can be obtained from the Expectation-Maximization (EM) algorithm (Dempster, Laird, & Rubin, 1977; Shumway & Stoffer, 1982). In the EM algorithm, the states \mathbf{u}_t are treated as missing values. Using previous estimates of $\boldsymbol{\mu}_0$, $\boldsymbol{\Sigma}_0$, $\boldsymbol{\Sigma}_u$ and σ_r , the missing states are imputed with their expected values (computed using the Kalman smoother), and new parameter estimates are then computed by maximising the joint likelihood of the parameters given the observations $R_{1:T}$ and (imputed) states $\mathbf{u}_{1:T}$. For further details, we refer the reader to Durbin and Koopman (2001) or Wu et al. (1996), which contains a computationally simpler implementation.

Appendix B

Assessment of the cues used

To determine the subjective cues (those on which participants actually based their predictions), we fitted DLMs with the following alternative cue vectors:

- Correct (C): $\mathbf{x}_t = (x_{1t}, x_{2t})$
- Interaction (I): $\mathbf{x}_t = (x_{1t}, x_{2t}, x_{1t} \times x_{2t})$
- Previous change (P): $\mathbf{x}_t = (x_{1t}, x_{2t}, y_{t-1})$
- Interaction + Previous change (I+P): $\mathbf{x}_t = (x_{1t}, x_{2t}, x_{1t} \times x_{2t}, y_{t-1})$

In addition to investigating which cues participants actually used, we also investigated the extent of individual differences in how participants adapted their cue utilization. Recall that the parameters of the DLM determine the initial utilization, the rate of adaptation, and the consistency of responses. By fixing these parameters to be identical for different participants, and comparing this to a model in which the parameters vary over individuals, we can investigate the extent of individual differences in the dynamic process of cue utilization. We fitted three versions of each model. In one version, we estimated the parameters separately for each individual participant, thus allowing the parameters to vary over individuals. In a second version, we estimated the parameters for each condition, thus assuming the parameter values to be identical within the conditions. In a final version, we assumed the parameter values were identical for all participants. The different restrictions on the model parameters result in a set of nested models for each of the four models (e.g., the version in which parameters vary between participants is nested under the version in which parameters vary between conditions). To compare nested models estimated by maximum likelihood, we can use likelihood-ratio tests to determine whether a particular restriction resulted in significantly poorer fit.

Experiment 1

The results of the likelihood tests comparing the models with identical parameter values within each condition to those with no equivalence between participants were $\chi^2(94) = 1022.67$, $\chi^2(164) = 1109.70$, $\chi^2(164) = 1109.37$ and $\chi^2(258) = 1186.46$, for models C, I, P and I+P respectively, all $p < .001$. These tests show that, for each model, the version in which parameters varied between participants fitted significantly better than the other versions. As such, there is sufficient evidence for individual variation. To compare the models fitted to the individual data, we used the Bayesian Information Criterion (BIC, Schwarz, 1978). This showed that model C fitted best for 23 participants, model I for 1 participant, and model P for 4 participants. The distribution of best fitting model was independent of condition, $p = .33$ (Fisher's exact test). Comparing the models fitted to the individual data showed that model C fitted best (the overall BIC values were 73984, 74620, 74545 and 75422 for models C, I, P and P+I respectively).

Experiment 2

The results of the likelihood ratio tests comparing the models with identical parameters within each condition to those with no equivalence between participants were $\chi^2(102) = 1488.69$, $\chi^2(178) = 1593.24$, $\chi^2(178) = 1585.00$ and $\chi^2(280) = 1704.12$, for models C, I, P and I+P respectively, all $p < .001$. Thus, the version with individually varying model parameters fitted significantly better than the other versions in which parameters varied between conditions or were constrained to be identical for all participants. A comparison between the models on the individual data showed that model C fitted best for 24 participants, model I for 2, and model P for 4 participants. The distribution of best fitting model was independent of condition, $p = .34$ (Fisher's exact test). Comparison of the models fitted to the individual data showed that model C fitted best (the overall BIC values were 87122, 87821, 87743 and 88698 for models C, I, P and I+P respectively).

Experiment 3

In fitting model P and model P+I, the value of the previous change in share price was set to 0 for trials in the Test Phase, as no outcome feedback was given there. The results of the

likelihood ratio tests comparing the models with identical parameters within each condition to those with no equivalence between participants were $\chi^2(142) = 1905.45$, $\chi^2(248) = 2088.42$, $\chi^2(248) = 2051.75$ and $\chi^2(390) = 2335.26$, for models C, I, P and I+P respectively, all $p < .001$. Thus, for each model, the version with individually varying parameters fitted better than the other versions. Allowing for individual variation in model parameters, model C fitted best for 26, model I for 3, and model P for 10 participants, while model I+P fitted best for one participant. The distribution of best fitting model was independent of condition, $p = .68$ (Fisher's exact test). Overall, model C fitted the data best (the overall BIC values were 84006, 84823, 84603 and 85699 for models C, I, P, and I+P respectively).

Appendix C Details of the learning models

This appendix contains details of the implementation of the five learning models. Free parameters in these models were estimated by maximum likelihood, using the Nelder-Mead simplex method as implemented in the “optim” function of the R environment for statistical computing (R Development Core Team, 2006). As for the DLMS, we fitted different versions of each model, assuming the parameter values to be identical for all participants (referred to as “././ind”), to be identical for participants within a conditions (referred to as “././con”), or to vary over individuals (referred to as “././ind”).

Bayesian Linear Filter (BLF)

The Bayesian Linear Filter is similar to the DLM used to estimate participants' cue utilization. Thus, in correspondence with the DLM, the Bayesian Filter makes the following assumptions

$$\begin{aligned} Y_t &= \mathbf{x}_t \mathbf{v}_t + e_t, & e_t &\sim N(0, \sigma_y^2), \\ \mathbf{v}_{t+1} &= \mathbf{v}_t + \mathbf{d}_t, & \mathbf{d}_t &\sim N(0, \Sigma_v), \\ \mathbf{v}_0 &\sim N(\boldsymbol{\mu}_0, \Sigma_0). \end{aligned} \quad (18)$$

Learning proceeds through the *filtering* distribution

$$p(\mathbf{v}_t | Y_{1:t}) = N(\mathbf{m}_t, \mathbf{C}_t), \quad (19)$$

the *state prediction* distribution

$$p(\mathbf{v}_{t+1} | Y_{1:t}) = N(\mathbf{m}_t, \mathbf{P}_t), \quad (20)$$

and the *forecast* distribution

$$p(Y_t | Y_{1:(t-1)}) = N(\mathbf{x}_t \mathbf{m}_{t-1}, V_t), \quad (21)$$

The parameters of these distributions are calculated as in Equations 8, 9, 11 and 12, replacing σ_r by σ_y and Σ_u by Σ_v .

Participants' responses are assumed to follow the forecast distribution, such that the distribution of a participant's response at time t is identical to the distribution in Equation 21. In fitting the model to participants' responses, the mean of the prior distribution was fixed to $\boldsymbol{\mu}_0 = (0, 0)^T$ and its covariance to the diagonal matrix $\Sigma_0 = \text{diag}(\sigma_0, \sigma_0)$. We set the covariance matrix of validity changes to the diagonal matrix $\Sigma_v = \text{diag}(\sigma, \sigma)$ in models “BLF/1/.”, and to $\Sigma_v = \text{diag}(\sigma_1, \sigma_2)$ in models “BLF/2/.”. The value of σ_y was freely estimated. In total, the Bayesian Linear Filter model had thus either three (σ_0 , σ , and σ_y) or four (σ_0 , σ_1 , σ_2 , and σ_y) free parameters.

Competitive Bayesian Filter (CBF)

The Competitive Bayesian Filter is based on the assumption that, at each trial, only a single cue is responsible for producing the criterion. Which cue is responsible can vary over trials. The criterion is assumed to be a linear function of the responsible cue. Thus, denoting the responsible cue on trial t as C_t , we can write the outcome on trial t , conditional upon j being the responsible cue, as

$$Y_t | (C_t = j) = v_{jt} x_{jt} + e_t \quad e_t \sim N(0, \sigma_y^2). \quad (22)$$

As in the BLE, the cue validities are assumed to vary over time as

$$v_{jt} = v_{j(t-1)} + d_{jt} \quad d_{jt} \sim N(0, \sigma_j^2), \quad (23)$$

starting from a prior distribution

$$p(v_{j0}) = N(\mu_{j0}, \sigma_{j0}^2) \quad (24)$$

In addition, the CBF incorporates an assumption regarding the changes in cue responsibility. More precisely, these are assumed to follow a (first-order) Markov process, such that the probability that cue j produces the criterion at trial t , given that k was the responsible cue on trial $t - 1$, is time-invariant. We will denote this conditional probability as

$$\beta_{jk} = P(C_t = j | C_{t-1} = k) \quad (25)$$

The Competitive Bayesian Filter combines the Kalman filter for learning cue validity with a discrete-state Markov model for inferring responsibility; such a model can be referred to as a Switching Kalman Filter (cf. Ghahramani & Hinton, 2000; Murphy, 1998). As in the BLE, learning of cue validity involves filtering distributions and the parameters of these are computed by the Kalman filter. However, in the CBF, a current observation only provides information of the validity of the responsible cue. Therefore, the observations are weighted according to the posterior probability of cue responsibility in the CBF. The posterior probability that cue k was responsible on trial t , $\alpha_{kt} = P(C_t = k | Y_{1:t}, \mathbf{x}_{1:t})$, can be computed as

$$\alpha_{kt} = \frac{\sum_{j=1}^m \alpha_{j(t-1)} \beta_{jk} N(Y_t | v_{kt} x_{kt}, \sigma_j^2)}{\sum_{k=1}^m \sum_{j=1}^m \alpha_{j(t-1)} \beta_{jk} N(Y_t | v_{kt} x_{kt}, \sigma_j^2)}, \quad (26)$$

where $N(a|b, c)$ denotes the Normal density with mean b and variance c evaluated at a . Weighting observations by α_{kt} is identical to rescaling the observational variance σ_y^2 as σ_y^2 / α_{kt} (Ghahramani & Hinton, 2000). Apart from this change, the computations are equivalent to those of the BLE, which can be simplified here by representing cue validity as a collection of univariate processes. Thus, if we let m_{jt} denote the mean of the filtering distribution of v_{jt} , then

$$m_{jt} = m_{j(t-1)} + K_{jt}(Y_t - x_{jt} m_{j(t-1)}), \quad (27)$$

where the Kalman gain is

$$K_{jt} = \frac{x_{jt} s_{jt}^2}{x_{jt}^2 s_{jt}^2 + \sigma_y^2 / \alpha_{jt}}, \quad (28)$$

and the variance of the state prediction distribution is

$$s_{j(t+1)}^2 = (1 - K_{jt} x_{jt}) s_{jt}^2 + \sigma_j^2. \quad (29)$$

In the CBF, the forecast distribution is a mixture of m Normal distributions (one component for each cue)

$$p(Y_t | Y_{1:(t-1)}) = \sum_{j=1}^m \pi_{jt} N(m_{j(t-1)} x_{jt}, V_{jt}), \quad (30)$$

where the mixing proportions are equal to the predicted responsibility

$$\pi_{jt} = P(C_t = j | Y_{1:(t-1)}, \mathbf{x}_{1:(t-1)}) = \sum_{k=1}^m \alpha_{k(t-1)} \beta_{jk} \quad (31)$$

and the variance of each Normal component is

$$V_{jt} = x_{jt}^2 s_{jt}^2. \quad (32)$$

The optimal point prediction for such a mixture is the weighted average

$$\hat{y}_t = \sum_{j=1}^m \pi_{jt} m_{j(t-1)} x_{jt}. \quad (33)$$

Participants' responses are assumed to be made as such a weighted average, but with a variance which reflects the uncertainty related to the individual components. More precisely, we assume participants' responses are distributed as

$$p(R_t | x_{1:t}, Y_{1:(t-1)}) = N(\hat{y}_t, \sum_{j=1}^m \pi_{jt}^2 V_{jt}), \quad (34)$$

In fitting the CBF, we assumed a symmetry in transitions between responsible cues, such that $\beta_{jk} = \beta_{kj} = \beta$ (note that, for $m = 2$, we have $\beta_{jj} = 1 - \beta_{jk}$). Furthermore, we fixed the mean of the prior distributions to $\mu_{j0} = 0$ with an equal variance for each cue, $\sigma_{j0}^2 = \sigma_0^2$. In versions "CBF/1/.", we assumed equivalence between the cues for the state variance $\sigma_j^2 = \sigma^2$, while in versions "CBF/2/.", this variance could differ between the cues. In total, the CBF had either 4 (β , σ_0^2 , σ^2 , and σ_y^2) or 5 (β , σ_0^2 , σ_1^2 , σ_2^2 , and σ_y^2) free parameters.

Least Mean Squares network (LMS)

The Least Mean Squares network is a simple single layer feedforward network with one input node for each cue and a single output node for the criterion. The activation of the input nodes is given by \mathbf{x}_t , and the activation of the output node is

$$\text{out}_t = \mathbf{x}_t \mathbf{w}_t. \quad (35)$$

The connection weights $\mathbf{w}_t = (w_{1t}, w_{2t})^\top$ are updated by the LMS or delta rule (e.g. Gluck & Bower, 1988)

$$\mathbf{w}_{t+1} = \mathbf{w}_t + \text{diag}(\boldsymbol{\eta})(y_t - \text{out}_t) \mathbf{x}_t^\top, \quad (36)$$

where $\boldsymbol{\eta}$ is a vector with learning rate parameters (such that each input-output connection can be learned at a different rate). The LMS network gives point predictions of participants' responses. To fit the model to participants' actual responses, we assumed that these were distributed as

$$p(R_t | Y_{1:(t-1)}, \mathbf{x}_{1:t}) = N(\text{out}_t, \sigma_r^2) \quad (37)$$

where σ_r^2 is constant variance which reflects the “implementation error”.

Starting weights were initialized to $\mathbf{w}_0 = (0, 0)^\top$. Two versions were fitted, one with identical, $\boldsymbol{\eta} = (\eta, \eta)$, and one with a different learning rate, $\boldsymbol{\eta} = (\eta_1, \eta_2)$, for the two cues. These are referred to as models “LMS/1/.” and “LMS/2/.” respectively. In total, the LMS model had thus either two (η and σ_r) or three (η_1 , η_2 and σ_r) free parameters.

Generalized Context Model (GCM)

In exemplar models such as the Generalized Context Model (Nosofsky, 1986), a response to a probe cue \mathbf{x}_t is made by recollecting previously encountered exemplars $\mathbf{x}_{1:(t-1)}$, and weighting the corresponding recollected criterion values $y_{1:(t-1)}$ according to the similarity of \mathbf{x}_t to \mathbf{x}_{t-k} , $k = 1, \dots, t - 1$. The similarity of two cues is a decreasing function of their distance

$$s_{t(t-k)} = s(\mathbf{x}_t, \mathbf{x}_{t-k}) = \exp(-\lambda(d_{t(t-k)})^q), \quad (38)$$

with $q = 1$ for the exponential, and $q = 2$ for the Gaussian similarity function. The distance is based on a generalized Minkowski distance

$$d_{t(t-k)} = d(\mathbf{x}_t, \mathbf{x}_{t-k}) = ((\mathbf{x}_t - \mathbf{x}_{t-k})^r \mathbf{w})^{(1/r)} \quad (39)$$

where $r > 0$ and $\mathbf{1}^\top \mathbf{w} = 1$ (i.e., the attention weights $w_i \geq 0$ sum to 1). To account for memory decay, exemplars were additionally weighted in accordance with their “distance” in time. At current time t , the time distance of an exemplar encountered at time $t' < t$ is $t - t' = k$ and the recency weight is

$$\pi_k = \exp(-\gamma k), \quad (40)$$

for exponential decay, or

$$\pi_k = k^{-\gamma}, \quad (41)$$

for power decay. Note that we have not included the usual linear scaling parameters in these functions, as these cancel out in Equation 42.

As in Juslin et al. (2003), we assume predictions are made as a weighted average of the criterion value of the stored exemplars

$$\hat{y}_t = \frac{\sum_{j=1}^{t-1} \pi_j s_{tj} y_j}{\sum_{k=1}^{t-1} \pi_k s_{tk}} \quad (42)$$

In addition, we assume responses are corrupted by Normally distributed random noise, so that the distribution of responses is

$$p(R_t | Y_{1:(t-1)}, \mathbf{x}_{1:t}) = N(\hat{y}_t, \sigma_r^2) \quad (43)$$

In the actual implementation of the model, we used the alternative parametrization

$$s_{t(t-k)} = \exp(-(\mathbf{v}^\top (\mathbf{x}_t - \mathbf{x}_{t-k})^r)^{q/r}), \quad (44)$$

where $\mathbf{v} = \lambda^{r/q} \mathbf{w}$, and $\mathbf{v} \in \mathfrak{R}^+$. The advantage is that the only constraint on \mathbf{v} is that the elements are greater or equal to zero. The values v_i represent the “stretching” (or contraction) of cue dimension i .

Parameters r and q were treated as fixed, and we fitted eight different versions of the GCM, resulting from the different combinations of $r = 1, 2$, $q = 1, 2$ and the two memory decay functions. The GCM has four remaining free parameters, $\mathbf{v} = (v_1, v_2)$, σ_r and γ . Tables D1 and D2 contain only the results from the models with $r = 2$ and $q = 1$, which fitted uniformly better than the other combinations of r and q (although the difference with the model with $r = 1$ and $q = 1$ was only slight).

Associative Learning Model (ALM)

The ALM is an associative network with one input node for each (possible) cue pattern, and an output node for each (possible) value of the criterion. We provide here a relatively straightforward extension of the ALM to multiple-cue functions, using multivariate Gaussian generalization gradients rather than the univariate Gaussian gradients of the original ALM. A probe cue pattern \mathbf{x}_t results in a distribution of activation across all input nodes, according to the generalization gradient

$$a_{it} = \exp\left(-.5(\mathbf{x}_t - \boldsymbol{\alpha}_i)\boldsymbol{\Sigma}^{-1}(\mathbf{x}_t - \boldsymbol{\alpha}_i)^\top\right) \quad (45)$$

in which $\boldsymbol{\alpha}_i = (\alpha_{1i}, \alpha_{2i})$ is the vector with locations for input node i , and $\boldsymbol{\Sigma} = \text{diag}(\sigma_1, \sigma_2)$ a scaling matrix. The activation of each output node is an additive function of the input activations

$$o_{jt} = \sum_i w_{ijt} a_{it}, \quad (46)$$

where w_{ijt} is the weight of the connection between input node i and output node j . These weights are initialized at 0, and updated according to the LMS rule

$$w_{ij(t+1)} = w_{ijt} + \eta(v_{jt} - o_{jt})a_{it}, \quad (47)$$

in which η is a learning rate parameter, and v_{jt} the training signal for output node j on trial t . This training signal is determined from the outcome feedback y_t by the Gaussian generalization gradient

$$v_{jt} = \exp\left(-.5(y_t - \mu_j)^2 / \sigma_y^2\right), \quad (48)$$

in which μ_j is the location of output node j and σ_y a scaling parameter.

Responses are a probabilistic function of the activation of the output nodes

$$P(R_t = \mu_j) = \frac{\exp(\lambda o_{jt})}{\sum_k \exp(\lambda o_{kt})}. \quad (49)$$

Note that in using the exponentiated version of Luce's ratio choice rule, our implementation differs slightly from that in Bussemeyer et al. (1997) and DeLosh et al. (1997), who used the simpler version $P(R_t = \mu_j) = o_{jt} / (\sum_k o_{kt})$. Using the exponentiated version was necessary, since there was no guarantee that the output values o_{jt} were positive. We also followed DeLosh et al. (1997) and used one input node for each unique presented cue pattern, rather than for each unique possible cue pattern. The number of input nodes was restricted for practical purposes, as there are 200^2 possible cue patterns, which would have resulted in a network with 200^3 connections. The restricted ALM is still rather large, with 60,000 input-output connections for each participant.

Appendix D
Fit of the learning models

This appendix contains the full results of the model fit for Experiments 1 and 2 (Table D1) and Experiment 3 (Table D2). The best fitting version of each model is signalled in bold.

Table D1: Fit measures of the learning models in Experiment 1 and 2.

model	Experiment 1			Experiment 2		
	# par.	BIC	R^2	# par.	BIC	R^2
BLF/1/all	3	78506	0.44	3	90750	0.24
BLF/2/all	4	78424	0.45	4	90757	0.24
BLF/1/con	6	78050	0.46	6	90700	0.24
BLF/2/con	8	77993	0.46	8	90605	0.25
BLF/1/ind	84	76319	0.56	90	88770	0.34
BLF/2/ind	112	75964	0.59	120	88736	0.36
CBF/1/all	4	78138	0.45	4	89422	0.32
CBF/2/all	5	78232	0.44	5	89412	0.32
CBF/1/con	8	77722	0.47	8	89350	0.31
CBF/2/con	10	77666	0.47	10	89352	0.31
CBF/1/ind	112	77124	0.54	120	88721	0.37
CBF/2/ind	140	76791	0.55	150	88716	0.37
LMS/1/all	2	78691	0.43	2	91193	0.21
LMS/2/all	3	78569	0.44	3	91067	0.22
LMS/1/con	4	78199	0.45	4	91149	0.22
LMS/2/con	6	78167	0.45	6	90939	0.23
LMS/1/ind	56	76689	0.55	60	89268	0.33
LMS/2/ind	84	76222	0.58	90	89328	0.34
GCM/exp/all	4	78831	0.42	4	90149	0.30
GCM/pow/all	4	79234	0.40	4	90504	0.27
GCM/exp/con	8	78523	0.43	8	90067	0.30
GCM/pow/con	8	78889	0.41	8	90412	0.28
GCM/exp/ind	112	77413	0.53	120	89209	0.38
GCM/pow/ind	112	77867	0.51	120	89675	0.35
ALM/1/all	4	77137	0.44	4	88359	0.28
ALM/2/all	5	77132	0.44	5	88353	0.29
ALM/1/con	8	76785	0.46	8	88355	0.29
ALM/2/con	10	76793	0.46	10	88349	0.29
ALM/1/ind	112	75479	0.53	120	87680	0.35
ALM/2/ind	140	74941	0.56	150	87691	0.36

Table D2: Fit of the learning models in Experiment 3

model	# par.	BIC					R^2				
		learning	test			all	learning	test			all
			inter.	extrap.	comb.			inter.	extrap.	comb.	
BLF/1/all	3	74117	3830	4603	8433	90863	0.18	0.33	-0.39	-0.15	0.15
BLF/2/all	4	73469	3861	4481	8342	90064	0.24	0.34	-0.36	-0.13	0.20
BLF/1/con	6	73004	3859	4497	8356	89661	0.28	0.34	-0.25	-0.06	0.25
BLF/2/con	8	72943	3855	4525	8380	89640	0.29	0.34	-0.28	-0.08	0.25
BLF/1/ind	120	70711	3709	4163	7871	86383	0.42	0.41	0.16	0.24	0.41
BLF/2/ind	160	70614	3710	4220	7930	86385	0.45	0.36	0.15	0.22	0.43
CBF/1/all	4	72368	3821	4245	8065	88550	0.30	0.28	0.04	0.12	0.29
CBF/2/all	5	72364	3820	4243	8062	88542	0.30	0.28	0.04	0.12	0.29
CBF/1/con	8	72335	3824	4294	8118	88611	0.31	0.27	0.01	0.10	0.29
CBF/2/con	10	72350	3822	4286	8108	88612	0.31	0.28	0.02	0.11	0.29
CBF/1/ind	160	71792	3739	4090	7829	87553	0.42	0.35	0.27	0.30	0.41
CBF/2/ind	200	71666	3732	4104	7836	87434	0.43	0.31	0.26	0.29	0.41
LMS/1/all	2	73687	3874	4412	8287	90198	0.22	0.36	-0.27	-0.06	0.20
LMS/2/all	3	73658	3874	4429	8304	90194	0.23	0.36	-0.29	-0.08	0.20
LMS/1/con	4	73614	3882	4420	8302	90186	0.23	0.33	-0.27	-0.07	0.20
LMS/2/con	6	73453	3878	4458	8336	90054	0.25	0.33	-0.32	-0.10	0.21
LMS/1/ind	56	71540	3764	4197	7961	87440	0.38	0.31	-0.02	0.09	0.35
LMS/2/ind	84	71416	3760	4172	7932	87264	0.42	0.29	0.05	0.13	0.38
GCM/exp/all	4	72968	3877	4182	8059	89198	0.30	0.27	0.14	0.19	0.29
GCM/pow/all	4	72953	3872	4216	8087	89183	0.30	0.28	0.10	0.16	0.29
GCM/exp/con	8	72654	3879	4247	8126	89016	0.33	0.26	0.08	0.14	0.31
GCM/pow/con	8	72589	3871	4282	8153	88935	0.33	0.28	0.04	0.12	0.31
GCM/exp/ind	160	71489	3840	4204	8044	87508	0.47	0.21	0.25	0.24	0.44
GCM/pow/ind	160	71660	3819	4440	8259	87889	0.46	0.27	0.15	0.20	0.43
ALM/1/all	4	71642	3850	4055	7904	87556	0.29	0.32	0.12	0.19	0.27
ALM/2/all	5	71599	3841	4049	7890	87490	0.29	0.32	0.12	0.19	0.27
ALM/1/con	8	71570	3870	4056	7926	87577	0.30	0.32	0.13	0.20	0.28
ALM/2/con	10	71488	3845	4046	7891	87403	0.29	0.33	0.12	0.19	0.28
ALM/1/ind	160	70286	3789	4121	7909	85997	0.41	0.32	0.00	0.11	0.38
ALM/2/ind	200	70090	3783	4322	8104	85940	0.44	0.29	-0.24	-0.06	0.39