Published in final edited form as:

J Exp Psychol Learn Mem Cogn. 2011 May ; 37(3): 673–687. doi:10.1037/a0022473.

Error Discounting in Probabilistic Category Learning

Stewart Craig,

School of Psychology, The University of Western Australia

Stephan Lewandowsky, and

School of Psychology, The University of Western Australia

Daniel R. Little

The University of Western Australia, Indiana University, and The University of Melbourne

Abstract

Some current theories of probabilistic categorization assume that people gradually attenuate their learning in response to unavoidable error. However, existing evidence for this error discounting is sparse and open to alternative interpretations. We report two probabilistic-categorization experiments that investigated error discounting by shifting feedback probabilities to new values after different amounts of training. In both experiments, responding gradually became less responsive to errors, and learning was slowed for some time after the feedback shift. Both results are indicative of error discounting. Quantitative modeling of the data revealed that adding a mechanism for error discounting significantly improved the fits of an exemplar-based and a rule-based associative learning model, as well as of a recency-based model of categorization. We conclude that error discounting is an important component of probabilistic learning.

Keywords

categorization; learning; error; relevance shifts

Throughout daily life, people make decisions based on uncertain cues that predict multiple outcomes: Doctors make diagnoses based on symptoms that are indicative of several potential diseases; dark clouds do not always bring rain, yet an overcast day is more likely to be associated with precipitation than sunshine; and despite injuries, trades, and home-field advantages, we learn to predict which teams are likely to win on the weekend. Accurate prediction requires learning probabilistic relationships between cues and uncertain outcomes (Kruschke & Johansen, 1999). Probabilistic categorization is one form of probability learning in which the to-be-predicted outcomes are discrete categorical values. Unlike its deterministic counterpart, in probabilistic categorization the same stimuli belong to multiple categories albeit with differing probabilities. In consequence, the same response to an identical stimulus can be reinforced as either correct or incorrect on a probabilistic basis (Kruschke & Johansen, 1999). This corrective feedback is often believed to drive category learning (Gluck & Bower, 1988a, 1988b; Rumelhart, Hinton, & Williams, 1986).

In deterministic category learning, people often perform sufficiently well to eliminate error completely. In probabilistic categorization, by contrast, avoiding errors and achieving perfect accuracy is by definition impossible. On that basis, Kruschke and Johansen (1999) proposed that during probabilistic categorization, people may eventually accept a certain

level of unavoidable error and progressively slow their learning, a phenomenon known as error discounting. Kruschke and Johansen (1999) instantiated error discounting within a computational model, called RASHNL (Rapid Attention SHifts 'N' Learning).

RASHNL includes an error discounting mechanism based on an “annealing” process, which works by reducing learning rates across time (e.g., Amari, 1967; Heskes & Kappen, 1991; Murata, Kawanabe, Ziehe, Müller, & Amari, 2002). This decline in learning rates causes a gradual decrease in the rate at which RASHNL updates associations. Thus, as learning progresses, errors are increasingly discounted because they contribute less and less to learning. The annealed learning provides the model with various advantages over a fixed learning rate: For example, annealing allows for large response shifts early in training, to quickly adapt to the environment, while the slowing of learning later in training permits the fine-tuning of probabilistic responding (e.g., Amari, 1967; Heskes & Kappen, 1991; Murata et al., 2002).

This article critically examines the notion of error discounting and seeks supporting evidence for its operation in probabilistic categorization. We start by examining existing evidence for discounting and find it to be surprisingly sparse. We next present two probabilistic categorization experiments in which the reinforcement probabilities for all stimuli were shifted during training, with the abruptness and timing of the shift varying between experiments and conditions. We present two forms of analysis to confirm the presence of error discounting. First, we examine decisional-recency effects in categorization. The analysis of decisional recency, as developed by Jones, Love, and Maddox (2006), pinpoints how people adjust their behavior in response to feedback from the immediately preceding trial. Using an extension of Jones et al. (2006)’s original analysis, we show that participants gradually come to ignore error feedback when making category decisions, and their sensitivity to feedback is only gradually restored after reinforcement probabilities change. Second, we applied four models of categorization to the results: An exemplar model without an associative-learning mechanism that bases decisions on a similarity comparison to all previously stored exemplars (the Generalized Context Model, GCM; Nosofsky, 1986), a model that utilizes only recency information (the Memory and Context model, MAC; Stewart, Brown, & Chater, 2002), a second exemplar model based on associative learning (RASHNL; Kruschke & Johansen, 1999), and a rule-based model that also included associative learning (derived from the hybrid Attention to Rules and Instances in a Unified Model, ATRIUM; Erickson & Kruschke, 1998, 2002). We find that only the associative-learning models provide an accurate account of the data. Additionally, the MAC, RASHNL, and ATRIUM models provide a significantly better account if they incorporate an error-discounting mechanism. We conclude that error discounting deserves to be given a prominent theoretical role in probabilistic categorization.

Error Discounting in Probabilistic Environments

Examination of the applied literature on performance in complex environments indicates that ignoring unreliable cues is a common risk in industrial settings (Lorenz & Parasuraman, 2007; Moray, 2007) and aviation (Wickens, 2007). However, the existing evidence for error discounting is equivocal. Empirical support for error discounting involves the finding that the delayed introduction of cues reduces their utilization (Edgell, 1983; Edgell & Morrissey, 1987). If cues suddenly become relevant after an initial period of non-relevance, people use them less than equivalent cues that were relevant from the outset. That is, people fail to notice a change in the environment and fail to learn (much) about a highly relevant but new cue (Kruschke & Johansen, 1999). Although this finding is suggestive of error discounting, other explanations, such as those for conventional blocking (Kamin, 1969; Rescorla & Wagner, 1972), cannot be ruled out. Blocking occurs when a previously learned association

between a cue, A, and an outcome, X, prevents learning of a new association between a second cue, B, and X, when trained in the presence of A (Kruschke & Blair, 2000). For example, after learning to associate a beep with an electric shock, further training with the conjunction of (the same) beep and a (novel) light (i.e., both the beep and the light now predict the ensuing shock) does not lead to any learning about the light. The light on its own will show little if any evidence of having been learned on a final test. Newly-relevant cues in probabilistic learning may be under-utilized because learning about the new cue may have been blocked by the association between the already-predictive cue(s) and the outcome. It is preferable to seek evidence for error discounting in situations in which there is only one stimulus dimension present: with a single cue, blocking cannot occur. The single-cue approach recognizes that if error discounting is present, people should be insensitive to shifts in reinforcement probabilities. For example, if a stimulus is initially placed in category A with 80% probability, and this suddenly shifts to 20% at some point during learning, error discounting implies that people might continue to predominantly assign that stimulus to category A, at least for some time.

Relevance Shifts in Probabilistic Environments

Research on the effects of shifts in reinforcement probabilities dates back several decades (Estes, 1984; Estes & Straughn, 1954; M. P. Friedman et al., 1964; Yelen & Yelen, 1969). However, upon closer inspection, those data turn out to be either ambivalent or of little relevance to the error-discounting notion. In a series of early studies (e.g., Estes & Straughn, 1954; M. P. Friedman et al., 1964; Yelen & Yelen, 1969) participants learned to predict which of two lights would be illuminated on the next trial. The actual probabilities with which the lights were illuminated differed from chance, and by associating the actual outcome (identity of the illuminated light) with their prediction, people's responses gradually came to track the actual event probabilities. Most relevant here is the fact that in all studies those probabilities changed during the experiment. For example, in the study by Yelen and Yelen (1969), one light might be illuminated on 90% of the first 100 trials (and the other light during the remaining 10%), with those probabilities suddenly reversing for the next 100 trials. In virtually all cases, participants were found to adapt to those changes with remarkable speed. At first glance, this outcome might seem to compromise the notion of error discounting. Upon closer inspection, however, the data arguably have little bearing on probabilistic categorization processes because people never associated distinct outcomes to different stimuli: The only "stimulus" in those early studies was the signal that denoted the beginning of a trial, and consequently any observed learning did not necessarily involve the formation of associations between stimuli and outcomes but simple awareness of base-rates of reinforcement. The data of Estes and Straughn (1954), M. P. Friedman et al. (1964), and Yelen and Yelen (1969) thus confirm that people remain sensitive to base-rates even after more than 1,000 trials with varying probabilities (M. P. Friedman et al., 1964); however, they tell us little about error discounting.

More recently, in a closer analogue to probabilistic categorization, Estes (1984) asked participants to choose between two alternatives on each trial. Each alternative, in turn, was associated with two outcomes (success or failure) that were reinforced with some varying probability. For one of the alternatives, the probability of success remained constant at .5 throughout. For the other alternative, the probability of success increased and decreased throughout training following a sine pattern. The results were intriguing: On the one hand, Estes found that participants were generally aware of the changing probability structure for one of the outcomes and responded accordingly (i.e., following the cyclical sine pattern, preferring that alternative whenever its probability of success exceeded .5 and rejecting it whenever it fell below .5). On the other hand, during a final transfer block involving a

uniform probability of success of .5 for both alternatives, people persisted with the cyclical pattern and showed little evidence of adaptation.

On balance, the existing data involving relevance shifts are, at best, ambivalent with respect to error discounting. The study that most closely resembled probabilistic categorization (Estes, 1984), simultaneously found evidence both for people's ability to track changing probabilities and against their ability to adapt after prolonged training. The latter outcome is suggestive of error discounting but remains inconclusive.

Present Experiments

The present experiments used a single, quasi-continuous cue that could take on four values; each value was probabilistically associated with two categories. The use of a single cue precluded the possibility of blocking. The principal manipulation in the experiments involved switching the reinforcement probabilities of the two categories part-way through training. Between experiments, we varied the rate at which the reinforcement probabilities were shifted, which was either sudden (Experiment 1) or gradual (Experiment 2). If the reinforcement probabilities are shifted gradually, the increase in error at any point during the shift is limited. That limited increase, in turn, might be more likely to be discounted. By contrast, a sudden shift in the reinforcement contingencies would lead to a large increase in the error signal, potentially alerting participants to the new environment. Within each experiment, we varied the point at which the shift occurred.

Experiment 1

The first experiment used a single sudden shift in reinforcement. Depending on condition, the shift occurred either early, near the middle, or late in the training sequence.

Method

Participants—Sixty-four members of the University of Western Australia campus community participated in the experiment. Participants were randomly assigned to one of three conditions; *early*, *mid*, or *late*. The number of participants in the three conditions was 21, 21, and 22, for the early, mid, and late conditions respectively. Participants received either course credit or A\$10 remuneration.

Stimuli and apparatus—The experiment was controlled by a Windows computer using a MATLAB program created with the aid of the Psychophysics toolbox (Brainard, 1997; Pelli, 1991). The task was a single-cue binary choice task that required participants to classify a set of 4 unique items into one of two categories, A or B. To control for potential idiosyncrasies of the stimuli, two different types of stimuli were used, either squares of varying size, or circles with a radial line at varying angles. Within each condition, participants were randomly assigned to stimulus type. For the squares, the 4 training items were of linearly increasing sizes of 44, 66, 88, and 110 mm edge length. The circles were all approximately 115 mm in diameter. The radial line within the circle was presented at increasing angles of 36, 72, 108, and 144 degrees (measured from the vertical). For some participants the radial line was presented in the left hemisphere of the circle whereas for others it was shown in the right hemisphere.

Stimuli were presented in black on a white background. The category membership of each stimulus was determined on a probabilistic basis, with the probability of an item belonging to category A at the outset of training equaling .2, .4, .6, and .8 in increasing order of size (or angle), respectively. Category B probabilities were $1 - P(A)$. Part-way through the experiment, those probabilities were shifted instantaneously, such that the probability of the

items belonging to category A changed to .8, .6, .4, and .2. Category B probabilities also shifted, and remained at $1 - P(A)$. The point at which this shift occurred varied between conditions: The shift occurred at the beginning of training block 2, 11, or 16 for the early, mid, and late conditions, respectively.

Procedure—Each trial commenced with a fixation symbol (“+”) in the center of the screen for 500 ms. The “+” was replaced by the stimulus, which was presented in the center of the screen along with the response options (i.e., Category A or Category B) underneath. The stimulus and response options remained visible until a response was made. Responses were made by pressing the “F” or “J” keys. Mappings of the two response keys to categories (A or B) were randomly determined for each participant. After each response, feedback (“CORRECT” or “WRONG”) was presented on the screen, below the item, for 1,300 ms.

Items were randomly presented in blocks of 40 trials, with each stimulus presented 10 times per block. Participants were trained for 18 blocks altogether. Self-timed breaks, lasting a minimum of 30 seconds, were presented after every four blocks.

Results and Discussion

In probabilistic categorization, above-chance performance can involve one of two primary behaviors, known as probability-matching and maximizing, respectively (Fantino & Esfandiari, 2002; Shanks, Tunney, & McCarthy, 2002). Probability matching is said to occur if people place items into particular categories with a frequency approximately equivalent to the probability of reinforcement: For example, if stimulus j is assigned to category A on 80% of all trials, then people probability match if they respond A with probability .8. By contrast, maximizing is said to occur if people always place an item into the category to which it is most likely to belong. Thus, for the above item j , maximizing implies that instead of responding with A 80% of the time, people would respond with A on all trials. In most situations, people tend to probability match rather than maximize (Fantino & Esfandiari, 2002; Shanks et al., 2002), although people’s performance may lie anywhere along that continuum.

To characterize performance, we computed probability-matching (PM) scores for the participants’ mean responses during the last pre-shift block(s); Block 1 for the early condition, Blocks 9 and 10 for the mid condition, and Blocks 14 and 15 for the late condition. PM scores were computed based on the following formula:

$$PM_{ij} = (R_i(A|j) - P(A|j)) * SI_j, \quad (1)$$

where $P(A|j)$ is the training probability for item j , and $R_i(A|j)$ is the probability of participant i responding with category A in response to item j . SI_j is a sign indicator for item j , such that $SI_j = -1$ if $P(A|j) < .5$ and $SI_j = 1$ if $P(A|j) > .5$ (D. Friedman & Massaro, 1998; Little & Lewandowsky, 2009b). Each participant’s PM score was based on the average score across the four items.

For ease of interpretation, the PM scores were then transformed to a scale ranging from 0 (chance responding) to 1 (full maximizing). For the present probability structure, a score of .4 indicates perfect probability matching. (Scores below 0 indicate reversed responding.) To ensure that only participants who learned the probability structures were included, participants with PM scores of less than 0 were removed from consideration ($N = 3, 4$, and 5 in the early, mid, and late conditions, respectively).

Trials with a response time (RT) of less than 150 ms or an RT more than 4 SDs above the participant's mean RT were removed from the analysis. Additionally, any participant who scored more than 10% of all RTs beyond these cutoffs was excluded from further analysis. Based on these criteria, an additional 3 participants were removed from each of the early, mid, and late conditions.

After the removal of all outliers, the mean PM scores for the remaining participants were 0.44 ($SD = 0.33$), 0.58 ($SD = 0.27$), and 0.50 ($SD = 0.26$), for the early, mid, and late conditions, respectively. These scores were significantly above chance (early: $t(14) = 5.21$, $p < .001$; mid: $t(13) = 8.12$, $p < .001$; late: $t(13) = 7.37$, $p < .001$).

To check whether there was an impact of stimulus type on responding, PM scores were next calculated for each of the 18 blocks. A three-way between-within subjects ANOVA (Condition $\times$ Stimulus type $\times$ Block), with the block PM scores as the dependent variable, uncovered no significant main effect of stimulus type, $F(1, 37) = 0.44$, $p > .1$. For the remainder of the analysis, the data were collapsed across stimulus type.

Mean response probabilities for each of the three conditions for each of the items across the 18 training blocks are presented in Figure 1. The figure makes two points: First, it shows that participants were able to adapt to the change in the probability structure in all three conditions. Second, it shows that people approximately probability-matched to all items, at least in the pre-shift blocks. For example, in the mid condition, the item that was reinforced as belonging to category A 80% of the time on pre-shift trials was placed in that category with probability $\approx .8$, and so on for all other items.

In order to obtain a more concise summary of people's adaptation to the shift, a single score was computed for each block that summarized responding to all four stimuli. By design, $P(A|j)$ across j was linear within a block (i.e., .2, .4, .6, and .8 or vice versa): Given that participants largely probability-matched, as evidenced by the mean PM scores, the slope of the line through $R_i(A|j)$ across j within each subject-block was likely also linear, a possibility confirmed by inspection of responses. The slope of the response probabilities within each block for each participant was determined by regression, and the observed mean response slopes and the slopes of the objective training probabilities for all blocks are presented in Figure 2. The figure shows each of the mean slopes with error bars indicating 95% confidence intervals. The confidence intervals provide statistical confirmation of a range of effects.

From the slope representation in Figure 2, we can observe participants' initial learning and subsequent adaptation to the shift in reinforcement probabilities. Initial learning was extremely rapid: from virtually the first training block onward, confidence intervals bracketed the objective slopes and were nowhere near chance (i.e., slope 0). This implies that most of the initial learning occurred during the first 40 trials.

Following the shift, participants crossed through chance responding by the second post-shift block in all conditions. However, people adapted more slowly to the change in probabilities than they did during initial learning. Inspection of individual responding suggested that the aggregate pattern was representative of behavior at the individual level. A visual indication of variation among individual responding is given by the 95% confidence intervals in Figure 2.

A two-way between-within subjects ANOVA examined whether the conditions differed with respect to the speed of adaptation. Condition (early vs. mid vs. late) was included as a between-subjects factor, training block was included as a within-subjects factor, with slope as the dependent variable. To render the conditions comparable, they were aligned around

the shift point, with the final pre-shift block and the two post-shift blocks comprising the levels of the block factor. There was a significant effect of block, $F(2, 80) = 28.74, p < .001$, $\eta_p^2 = .42$, but no significant effect of condition, $F(2, 40) = 0.29, p > .1$, and no interaction, $F(4, 80) = .32, p > .1$. Thus, the timing of the shift in the training sequence did not significantly affect the rate of adaptation.

We next sought direct evidence for error discounting by examining trial-by-trial response contingencies. Jones et al. (2006) introduced a trial-by-trial analysis technique to detect “decisional-recency” effects in probabilistic categorization. Specifically, Jones et al. showed that when the stimulus on the current trial was identical to the immediately preceding one, people were biased toward giving the same category response as the feedback presented on the previous trial. As the perceptual difference between the current and the previous stimulus increased, people’s response on the current trial was increasingly more likely to *differ* from the preceding category feedback. This decisional recency is consistent with the notion that people’s responding is, at least in part, driven by the perceptual contrast between successive stimuli (cf. Stewart et al., 2002). In a nutshell, when a stimulus was repeated, people responded with the category just reinforced whereas they shifted to the alternate category when a stimulus differed dramatically from its immediate predecessor. (Intermediate stimulus differences caused intermediate tendencies to shift.)

Here, we extend Jones et al. (2006)’s technique to explore the existence of error discounting. Whereas Jones et al. focused on the impact of category feedback—irrespective of whether the response was correct or not—on responses to the following stimulus, we examined whether people’s response on trial n was affected by whether or not their response on trial $n - 1$ was *correct*. That is, whereas Jones et al. considered only the nature of the feedback on trial $n - 1$, we additionally determined whether or not the probabilistic feedback $n - 1$ was identical to the response $n - 1$.

If people’s responding is subject to decisional recency, then the following contingencies should be observed. If the same stimulus is repeated on two successive trials, people should repeat their response if it was correct on the first trial but they should shift to the alternate response following an error. Conversely, if the stimuli on two successive trials are maximally different, then an error should be followed by the *same* response whereas a correct response should be followed by a shift. (Intermediate stimulus differences should again give rise to intermediate response tendencies.) If people gradually come to discount errors, then those response contingencies should decrease across training at least up to the shift in reinforcement probabilities. After the shift, those response contingencies might re-emerge, albeit perhaps slowly.

Figure 3 shows the trial-by-trial contingencies for the mid and late conditions combined across training blocks. (The early condition was omitted because it only comprised a single pre-shift training block.) The top panel shows the probability with which people shifted their response from one trial to the next, $P(\text{shift})$, for trial pairs involving the same stimulus. The bottom panel shows $P(\text{shift})$ for maximally different stimuli. $P(\text{shift})$ was computed by aggregating across trials and participants simultaneously. In order to render the conditions comparable, the data were “end-point aligned,” such that the first 5 blocks of the mid condition were aligned with the first 5 blocks of the late condition. Likewise, the final 5 pre-shift blocks in the mid condition, Blocks 6 to 10, were aligned with the last 5 pre-shift blocks in the late condition, viz. Blocks 11 to 15. After alignment, block numbers were zero-centered on the last pre-shift block.

The figure permits several conclusions: First, there is clear evidence for decisional recency because for an identical stimulus (top panel), people tend to shift to the alternate category

following an error but they repeat the same category after a correct response. Those response tendencies are reversed for the maximally different stimuli (bottom panel). Second, there is evidence for error discounting because during the pre-shift blocks, $P(\text{shift})$ gradually decreases following an error when the items are the same, and increases for items that are different. That gradual process is reversed once reinforcement probabilities are experimentally shifted and people become more sensitive to errors. We provide statistical confirmation of this pattern in a joint analysis of both experiments once all data have been presented.

Experiment 2

In Experiment 1, we focused on the consequences of sudden shifts in the probabilistic environment. The purpose of Experiment 2 was to explore how people respond to more gradual changes in the probabilistic structure. The consequences of error discounting may be more pronounced and more prolonged when the probabilistic environment shifts more gradually.

Method

Participants—The participants were 38 members of the University of Western Australia community. Participants were randomly assigned to the two conditions, early ($n = 18$) and late ($n = 20$). Of these participants, 9 were presented with square stimuli, and 29 were presented with the circle and radial line stimuli. Participants received either course credit or A\$10 remuneration.

Stimuli and apparatus—The stimuli and apparatus were almost identical to that in Experiment 1, with the following exceptions. As in Experiment 1, the feedback was given on a probabilistic basis. Initially, the probability of an item belonging to category A was .2, .4, .6, and .8, respectively, in increasing order of square size or radial line position. Category B probabilities were $1 - P(A)$. In the early condition, the probabilities shifted from their starting values at the beginning of Block 7, viz. to .35, .45, .55, and .65, respectively, for the four items. At the beginning of Block 9, the probabilities shifted through chance to .65, .55, .45, and .35, before finally settling at .8, .6, .4, and .2 from the beginning of Block 11. The probability of an item being in category B was always equal to $1 - P(A)$. In the late condition, the probabilities shifted under the same schedule, except that the shift occurred later in the sequence, beginning with Block 11 and finishing with Block 15. The procedure was identical to Experiment 1.

Results and Discussion

PM scores were calculated over Blocks 5 and 6 for the Early condition, and 9 and 10 for the Late condition. Two and five participants were excluded from the early and late conditions respectively, as their PM scores were below chance. As in Experiment 1, trials with an RT of less than 150 ms or an RT of more than 4 SD above a participant's mean RT were removed from further analysis. Participants were removed if more than 10% of their responses exceeded the RT criterions. Based on this criterion, an additional four participants were removed from the early condition, and five from the late condition. The mean PM scores for the remaining participants in the early and late conditions were 0.64 ($SD = 0.27$) and 0.65 ($SD = 0.22$). The mean scores for the early ($t(11) = 8.18, p < .001$) and late ($t(9) = 9.33, p < .001$) conditions were significantly greater than chance. One additional participant was removed from the late condition for the slope plots and for the stimulus check. After removing trials with extreme reactions times, this participant had no responses for two items in the final block, thus response slopes and PM scores could not be calculated for the final block. This participant was retained for all other analyses (unless otherwise specified).

As in Experiment 1, PM scores were next calculated for each of the 18 blocks. A three-way between-within ANOVA (Condition $\times$ Stimulus type $\times$ Block) was run with these PM scores as the dependent variable. The ANOVA uncovered no significant main effect of stimulus type, $F(1, 17) = 0.20, p > .1$. Data were collapsed across stimulus type for the remaining analysis.

Response slopes were computed in the same manner as for Experiment 1, and are shown in Figure 4. In both conditions, participants adapted to the shift. There was, however, notably little change in responding after the first partial shift in training probabilities (i.e., in Blocks 7 and 8 in the early condition and in Blocks 11 and 12 in the late condition). Participants appeared to adapt more to the second shift, with mean slopes crossing through chance responding by Blocks 12 and 15 for the early and late conditions respectively.

Participants appeared to learn more slowly after the shift than they did during initial learning. The contrast in learning speeds between the pre- and post-shift learning was more evident than in Experiment 1. In both conditions, pre-shift response slopes were above the objective response slopes by Block 1. However, after the shift, the mean response slopes do reach the objective training probabilities in either conditions.

A two-way between-within subjects ANOVA was conducted on the slope estimates. The between-subjects factor was condition (early vs. late). The within-subjects factor was blocks. Due to the gradual-shift design of the present experiment, the training-block factor had 8 levels: these ranged from the two final pre-shift blocks to the second block following the final shift. Thus, for the early condition, the 8 blocks were Blocks 4 to 12, and for the late condition, they were Blocks 8 to 16. A Huynh-Feldt correction was used to correct for violations of sphericity. The change in slope across blocks was significant, $F(3.72, 74.39) = 23.35, p < .001, \eta_p^2 = .54$, but there was no significant effect of condition, $F(1, 20) = 0.57, p > .1$, and no interaction, $F(3.72, 74.39) = 0.75, p > .1$.

A decisional-recency analysis was conducted as in Experiment 1. The results of this analysis are shown in Figure 5, with the upper and lower panels showing the $P(\text{shift})$ for identical and maximally different pairs of stimuli, respectively. As for the first study, data were simultaneously aligned on the first block and the block at which the training probabilities shifted; that is, Block 6 for the early condition and Block 10 for the late condition. As in Experiment 1, the $P(\text{shift})$ can be seen to decrease for identical items and increase for the maximally dissimilar items with training.

Statistical confirmation of the apparent error discounting revealed by the decisional-recency analysis in both Experiments 1 and 2 was provided by a multilevel generalized linear model (Baayen, Davidson, & Bates, 2008). To maximize power of the analysis, the mid and late conditions of Experiment 1 were analyzed together with the mid and late condition from Experiment 2. The early condition of Experiment 1 was excluded from the analysis as it had only one pre-shift block. Because the decisional-recency analysis pointed to the gradual discounting of errors *before* reinforcement probabilities were shifted, only pre-shift blocks were considered for each of the conditions. Thus, from Experiment 1, the analysis included Blocks 1 to 10 and 1 to 15 for the mid and late conditions respectively. From Experiment 2, the analysis included Blocks 1 to 6 and 1 to 10 for the early and late conditions respectively. To minimize the influence of the less diagnostic trial-pair transitions involving moderately dissimilar pairs of stimuli, only perceptual differences 0 (identity) and 3 (maximally dissimilar) were considered. The dependent variable was whether or not the response on the current trial differed from the preceding response (coded as 1 or 0). The perceptual difference between the two consecutive stimuli (0 or 3) was entered as a fixed-effect factor

and was fully crossed with block and with a factor called ‘error-before’ which captured whether or not an error had been made on the previous trial.

The results are summarized in Table 1 and confirm the pattern in Figures 3 and 5. We highlight four effects: (a) The greater the perceptual difference between two successive stimuli, the more likely people were to switch their response category. (b) An error was more likely to be followed by a switch than a correct response. (c) Those two variables interacted such that a switch became less likely with increasing perceptual difference if an error had just been committed (compare the two panels of Figure 3 for a particularly clear illustration of this interaction). (d) Crucially, the role of an error declined with blocks ($EB \times$ block interaction), thus attesting to the presence of error discounting.

In summary, the trial-by-trial contingency data from both experiments replicated the decisional recency previously reported by Jones et al. (2006). We additionally established the presence of error discounting: People gradually adjusted the probability with which they shifted their response on trials following errors.

Computational Modeling

Both experiments provided multiple indications for the presence of error discounting. First, there was evidence for slower post-shift learning than during the initial learning; this was particularly evident in Experiment 2. Second, in Experiment 2, there was a notable delay in adaptation to the new reinforcement probabilities. Third, the decisional-recency analysis showed that at a trial-by-trial level, people gradually came to ignore errors as training progressed, even before reinforcement probabilities were shifted. We next turn to quantitative modeling to investigate further the existence of error discounting. The aim of using quantitative modeling was twofold. First, to verify whether an error-discounting mechanism could improve the ability of computational models to account for the data. Second, to pursue a theoretical exploration of the processes underlying error discounting.

We applied four models to the data. One model, the GCM (Nosofsky, 1986), was used to explore a potential alternative explanation for error discounting based on sample size. The other three models were the MAC model (Stewart et al., 2002), RASHNL (Kruschke & Johansen, 1999), and ATRIUM (Erickson & Kruschke, 1998). In each of the latter three models we instantiated a computational error-discounting mechanism. The aim of including an error-discounting mechanism was to determine whether, after controlling for the number of parameters, theories of categorization with an error-discounting mechanism were able to better account for the data than theories without error discounting. We give a brief summary of each model and the modeling results below. Full details of the models, model fitting procedures, and model performance can be found in the web supplement.

GCM

The GCM is an exemplar-based model that stores representations of all previously encountered stimuli. The model categorizes new stimuli by making a similarity comparison between the new stimulus and all stored exemplar representations. The GCM was used to explore an alternative explanation for error discounting based on sample size. It is possible that shifts in the feedback probabilities might have a small impact on behavior simply due to the large number of stored exemplars consistent with the initial training probabilities which “drown out” the later (and fewer) exemplars stored with the revised training probabilities. Due to its exemplar-based architecture, the GCM was a fitting choice to explore this sample-size explanation. The GCM contains two free parameters. The first is a specificity parameter, c . The specificity parameter controls the rate at which the psychological similarity to a stored exemplar decreases with increasing perceptual distance from a given

stimulus. The second parameter is a scaling parameter, γ , that controls the decisional criterion of the model. The γ parameter can control whether the model responds at chance, probability matches, or maximizes.

MAC

Unlike most of its rivals, the MAC model uses only feedback and perceptual information from the previous trial (or limited set of trials) to make a category decisions. If the previous item was categorized correctly, then the MAC is more likely to give the same response to a perceptually similar item on the following trial. Conversely, the model is more likely to give an alternative response to a perceptually different item. This utilization of information from the immediately preceding trial(s) only makes the MAC model a prime candidate for exploring the impact of error discounting at the level of trial-by-trial decisional recency. Here we use the original MAC (Stewart et al., 2002), as opposed to an extended version of the model (Stewart & Brown, 2004). As noted by Stewart and Brown (2004), with greater trial memory, the MAC approaches the behavior of the GCM. Thus, the original MAC (with a single-trial memory) can provide greater comparative contrast to the GCM than can the extended MAC.

The MAC has a single parameter, b , which controls the rate at which the psychological similarity to the previous stimuli decreases with greater perceptual similarity between the current and previous items. We instantiated an error-discounting mechanism in the MAC by gradually attenuating the role of the previous stimulus for trials following an error. Specifically, we used two separate values of b following correct and following incorrect responses, respectively, and the value of b following incorrect responses was reduced by a factor r on each trial, given by:

$$r = 1/(1+\rho), \quad (2)$$

where ρ was a free parameter that controlled the rate at which the similarity parameter, b , was reduced.

RASHNL

RASHNL is an exemplar-based model of associative learning. The model is an extension of the exemplar model ALCOVE (Kruschke, 1992). Like the GCM, RASHNL makes categorization decisions on the basis of a similarity comparison between the present stimulus and all stored exemplar stimuli. Unlike the GCM, RASHNL includes an error-based, associative-learning mechanism that controls the strength of association weights between exemplars and categories. The standard version of RASHNL can also learn to distribute its attention across stimulus dimensions. Due to the use of a single-dimension task, the attention features of RASHNL were not required for the present simulations.

RASHNL was chosen for its built-in error-discounting mechanism. The mechanism in RASHNL works by decreasing, or “annealing,” the rate at which the model learns the associations between exemplars and categories. On each trial, an association learning rate, λ , is multiplied by the factor r , given in (2). As before, ρ is a free parameter controlling the rate at which the learning rate is reduced across trials. Higher value of ρ lead to fast annealing. As RASHNL uses errors to learn, by annealing the learning rates, as training progresses, RASHNL effectively ignores errors. RASHNL has two additional parameters, c , and ϕ , that are equivalent to the GCM’s specificity parameter, c , and the decision criterion parameter, γ , as outlined above.

rATRIUM

To extend the potential generality of our conclusions beyond exemplar- and recency-based models, we also included a rule-based variant of ATRIUM (Erickson & Kruschke, 1998). At present both rule and exemplar models are popular within the probabilistic categorization literature, with prior research providing conflicting support for each type of model (Juslin, Jones, Olsson, & Winman, 2003; Kalish & Kruschke, 1997; McKinley & Nosofsky, 1995; Rouder & Ratcliff, 2004). As the present study is not designed to distinguish between these models types, we included a rule-based model to ensure that any results were not specific to a particular type of theory. The full version of ATRIUM is a hybrid rule and exemplar model. For present purposes we used only the rule module, which we term, *r*ATRIUM. *r*ATRIUM categorizes items by dividing the category space by a rule boundary. Each side of this boundary is associated with a particular category. The further a new item falls on one side of the boundary, the more likely it is that the item will be placed in the associated category. Like RASHNL, *r*ATRIUM has an associative-learning mechanism. This associative learning mechanism allows *r*ATRIUM to learn associations between categories and each side of the category boundary. To explore error discounting, we implemented the same error-discounting mechanism in *r*ATRIUM as in RASHNL. As in RASHNL, *r*ATRIUM's learning-rate parameter, λ , was annealed across training. The parameter, λ , controls the rate at which *r*ATRIUM learns associations between categories and each side of the category boundary. The annealing rate was controlled by the free parameter ρ , as given in (2). *r*ATRIUM includes two additional free parameters, a noise parameter, μ , which controls the standard deviation of perceptual noise around the rule boundary. Greater values of μ , decrease the ability of the model to distinguish an item's location relative to the rule boundary. The final parameter, ϕ , is a decision criterion parameter, equivalent to the γ and ϕ parameters in the GCM and RASHNL, respectively.

Model Fitting Procedures and Modeling Results

For the three models with an error-discounting mechanism, the MAC, RASHNL, and *r*ATRIUM, we ran two sets of fits, one with the error-discounting mechanism functioning, and one with it switched off. The models were fit at the level of individuals by maximizing the binomial log-likelihood.

Table 2 summarizes the fit statistics (i.e., negative log-likelihoods, Akaike information criterion, AIC_c , and Bayesian information criterion, BIC) for each of the models (figures of the model predictions can be found in the web supplement). From the table, the GCM provided a poor account of the results, with relatively high AIC_c and BIC values than compared to RASHNL and *r*ATRIUM. From inspection of the model predictions, the GCM was unable to adapt to the shift as fast as participants did in all conditions, with the exception of the early condition of Experiment 1. The poor performance of GCM rules out the sample-size explanation of error discounting.

RASHNL, with annealing on, provided the best account of the data. Inspection of the model predictions showed that RASHNL provided a good quantitative fit of the data, closely capturing the participant's behavior throughout the shift in all conditions. The fits of RASHNL with annealing on were followed by *r*ATRIUM with annealing on, and then RASHNL and *r*ATRIUM with annealing off. The MAC did not provide a good quantitative fit; however, inspection of the best-fitting predictions of the model showed that the MAC qualitatively captured the data, including the shift in reinforcement.

Crucially, for all three models with an error-discounting mechanism, the fits were better when the mechanism was turned on than when it was off. For the MAC, RASHNL, and *r*ATRIUM, lower AIC_c and BIC values were obtained when annealing was on than when

annealing was off. Likelihood-ratio tests confirmed that each of the models fit significantly better when they included an annealing mechanism (RASHNL: $\chi^2(65) = 1034.81, p < .001$, *r*ATRIUM: $\chi^2(65) = 1313.56, p < .001$, MAC: $\chi^2(64) = 2547.14, p < .001$). The fact that all three models fit better with error discounting suggests that (a) error discounting is required to accurately capture probabilistic categorization behavior, and (b) that this holds true regardless of the particular type of category representation.

General Discussion

Summary of Findings

The present article sought to investigate (a) whether people gradually tend to discount errors during probabilistic categorization, and (b) what the underlying mechanism of error discounting are, if present. Taken together, the results clearly reveal the presence of error discounting during probabilistic category learning. Error discounting can be inferred from (a) the difference between the initial fast learning and the slower post-shift learning, (b) the apparent lack of adaptation to the first partial shift in Experiment 2, (c) the decline in likelihood with which participants shifted their category response on trials following an error, and (d) the fact that both associative-learning models and the MAC model handled the data significantly better when they included an error-discounting mechanism than when they did not.

The present evidence for error discounting is generally consistent with previous findings (Edgell, 1983; Edgell & Morrissey, 1987) that implicated the presence of annealing in models of categorization (Kruschke & Johansen, 1999). Unlike those precedents, however, the unidimensional design of the present studies precluded alternative interpretations based on blocking. Before we discuss the implications of our results, we briefly address some possible limitations.

Limitations and Concerns

Participants in our studies were not informed about a possible change in the probabilistic environment. In the natural world, by contrast, people may be more wary of changing probabilities and might thus adapt faster than observed here. However, the fact that learning in Experiment 2 continued to be slow after participants had transited through chance responding—at which point they would likely be aware of the changed reinforcement probabilities—suggests that awareness of a shift might not prevent (or undo) error discounting. Moreover, the decisional-recency analysis that showed a decline in the probability of shifting responses following errors, was run on pre-shift blocks only. Thus, people were shown to discount errors even before any shift had taken place. The shift in reinforcement probabilities thus helped highlight the existence of error discounting; however, it was not critical for error discounting to occur and to be detected.

Critics might also puzzle why the extent of pre-shift training failed to have any effect. If error discounting is gradual, why did people's adaptation rate not decrease with more pre-shift training? Our response is twofold: First, given the relative simplicity of our task, even limited training may have been sufficient to induce at least some discounting. This possibility is supported by the finding that even after the very first block of training, performance was far above chance and probability matching close to the objective target. It is also supported by the visible trends in the decisional-recency analysis. Second, RASHNL and *r*ATRIUM required annealing for their account of the data, irrespective of the extent of pre-shift training and even when parameter values were constrained to be identical across conditions. This finding confirms that error discounting commenced early on during training and then continued throughout the experiment.

Turning to potential theoretical limitations, we note that the use of a single stimulus dimension—while eliminating an alternative blocking explanation—rendered the attention and gain features of RASHNL and *r*ATRIUM irrelevant. As it is assumed in RASHNL that the attention and gain learning rates are also annealed (Kruschke & Johansen, 1999), the role of these features must therefore await further exploration. Likewise, it remains unclear how discounting operates in multi-dimensional environments. On the one hand, error discounting may play a greater role in complex environments, because if people have more difficulty determining the underlying probability structure, they may show a greater tendency to ignore errors. On the other hand, it may be that with multi-dimensional stimuli, error discounting is overshadowed by blocking effects, and thus has little impact on its own.

Finally, error discounting might interact with other effects and strategies present in multi-dimensional probabilistic categorization. For instance, Little and Lewandowsky (2009a) showed that compared to a deterministic condition, participants learning an identical category structure under probabilistic reinforcement spread their attention across a wider selection of possible cues. Specifically, people detected a non-diagnostic correlation among cues that escaped notice under deterministic learning conditions. It follows that in multi-dimensional categorization, error discounting might only have an impact once alternative strategies (i.e., attentional allocations) have been exhausted. Those issues await resolution by future research.

Theoretical Implications

In the present study we used three types of categorization models; an exemplar model, a rule-based model, and a recency-based model, to explore the existence of error discounting. The present study was not designed to distinguish between these types of categorization models, nor was it designed to determine how people represent categories. Instead, the inclusion of the three model types demonstrated that regardless of how categories are represented, error discounting is important to account for probabilistic categorization behavior. We therefore conclude that annealing of learning rates and error discounting deserve a prominent theoretical role in approaches to probabilistic categorization.

If we accept the necessity for error discounting, how is it best described at a conceptual level? Does error discounting reflect an on-going and ever-increasing habituation to a constant error signal? Alternatively, does error discounting reach a plateau, thus effectively differentiating between two error signals—one that is unavoidable due to the probabilistic environment and (potentially) another one that is unexpected and arises from changes in the environment?

The former alternative corresponds to the instantiation of annealing in RASHNL and *r*ATRIUM, whose learning rates ultimately and inevitably reach 0 (whereupon no further adaptation to a shift occurs, as confirmed by simulations involving a large amount of additional pre-shift training that are not reported in detail here). If people, as assumed by the models, progressively discount more and more error, how would they notice a gradual shift in reinforcement probabilities (as in our Experiment 2), and how would they learn from it? One possibility is that not all feedback-based learning is error driven. Indeed, there is recent neuroscientific (Histed, Pasupathy, & Miller, 2009) and behavioral (Smith & Kimball, 2010) evidence that people are capable of learning from feedback on correct responses. These recent findings run contrary to the common error-focused assumptions of connectionist models (e.g., Gluck & Bower, 1988a, 1988b; Rumelhart et al., 1986). We are not aware of any connectionist learning that is driven by *correct* responses, although these recent results (Histed et al., 2009; Smith & Kimball, 2010) may stimulate developments in that direction.

The latter alternative, namely that error-discounting reaches a plateau, is not currently instantiated in RASHNL and rATRIUM but nonetheless seems rather attractive for two reasons: First, people in both experiments *were* responsive to shifts, albeit sometimes with a considerable delay (in Experiment 2). Second, the extent of pre-shift training did not affect people's responsiveness to the shift in either experiment. This result is difficult to reconcile with a notion of growing habituation to error and it is more compatible with the idea that a *certain amount* of error is discounted, but that any further error is used for associative learning.

The idea that people discount some variability in the error signal resonates with the idea that people adjust their strategies avoid *over-fitting* or, in other words, overcompensating for noise (Myung & Pitt, 2004). The idea of over-fitting is germane if we assume that participants are in effect trying to distinguish some true feedback signal from variation or noise within that signal. Hence, participants might be aware that the overall environment has changed but are unaware of the extent of that change. Error discounting can therefore be viewed as a rational way to exclude excess variation from influencing learning. This is similar to modern machine learning algorithms which typically include a regularizing term (i.e., a penalty on the fitting measurement) to effectively limit the types of patterns that can be learned and reduce over-fitting (Bishop, 2006). Other approaches, such as the Bayesian learning approach we consider below (see e.g., Bishop, 2006, Chapter 3), implement regularization through use of prior probability distributions over relevant patterns.

The observation that error discounting exists at the level of decisional-recency flags the importance of considering corrective feedback when exploring sequential effects in categorization. Previous analyses of decisional-recency effects (e.g., Jones et al., 2006; Stewart et al., 2002) have focused on the impact of category-label *feedback* irrespective of whether or not the previous response was an error. As we have observed, people react differently following errors and correct responses; additionally, people adjust their behavior in response to errors over time. These reactions to errors in categorization suggest that it is important to consider corrective feedback when exploring decisional-recency effects.

So far, the focus here has been on participants' tendency to ignore error-based corrective feedback information. While participants did tend to discount errors, and thus ignored feedback, participants simultaneously showed signs of precise utilization of feedback information, as evidenced by people's highly-accurate probability-matching behavior. In the present tasks, the probability with which people categorized items approximated the probability with which the items belonged to each category. This accurate probability matching cannot arise without accurately observing the feedback information. Thus, participants simultaneously exhibited precise utilization of feedback information and a tendency to ignore feedback information. This duality of the utilization of feedback information is counter-intuitive: How can people dismiss feedback while being demonstrably attuned to it at the same time? The answer is implied by our modeling: The probability-matching behavior is established early on, when associations between stimuli and outcomes are learned on the basis of feedback. Once established, those associations persist (unless a reversal shift takes place) despite the increasing discounting of errors and the associated feedback.

Links to Bayesian Perspectives on Category Learning

The recent popularity of probabilistic-descriptive approaches to cognition (e.g., Chater, Tenenbaum, & Yuille, 2006) raises the possibility that error discounting might be tractable within a Bayesian approach. Bayesian approaches to cognition take into account our beliefs and expectations in the form of a prior probability distribution over potential hypotheses. Our prior beliefs are transformed into *posterior* beliefs in light of new data; hypotheses

which are consistent with the data are maintained whereas inconsistent hypotheses are discarded. In our task, as knowledge accumulates about the probability with which each item falls into each category, the distribution that represents those beliefs becomes more peaked and less variable with continued training. Hence, the influence of each individual stimulus on the posterior expectations of an item's category membership decreases: Less variable prior distributions require more new evidence (i.e., higher error rates) in order for new hypotheses to have any weight in the posterior. Error discounting may thus be interpretable in Bayesian terms as reflecting the continued "sharpening" of the distribution over the reinforcement rate of each item.

Implications of Error Discounting

Although our work was confined to probabilistic categorization, we can draw some connections to other situations involving decision making and industrial processes.

One perspective on error-discounting is that people minimize cognitive "expenditure" by dismissing the feedback signal that would otherwise entail further learning. This reduction in cognitive expenditure is also at the heart of the simple heuristics that underlie much of human decision making, and which have been characterized as maximizing our return on the investment of cognitive effort (Gigerenzer, Todd, & ABC Research Group, 1999). Similar to decision making heuristics—such as the recognition heuristic (Goldstein & Gigerenzer, 1999, 2002)—error discounting is highly adaptive for environments with unavoidable error, in which continuing to attend to errors entails no benefit for accuracy, whereas ignoring errors may free up limited cognitive resources.

However, in the same way that decision-making heuristics may incur a cost when the environment fails to conform to people's expectations (Tversky & Kahneman, 1974), error discounting can also incur a high cost in some environments. For example, pilots who ignore frequent warnings from an autopilot system are at increased risk of causing an accident (Farrell & Lewandowsky, 2000; Parasuraman & Riley, 1997; see also, Moray 2007). High false alarm rates for automated error detection systems result in people slowing their response to, or simply ignoring warnings about potential errors (Parasuraman & Riley, 1997). For example, in a task monitoring the navigation, speed, power and cargo systems of a ship, if the system sets off frequent false error alarms, participants eventually ignore these warnings signals (Kerstholt & Passenier, 2000). With frequent false alarms, participants made fewer attempts to gain information needed to assess the problem, and were slower to access this information when they did. This behavior constitutes a good analog to the error discounting observed in our studies. However, in many real-life situations, poor reliability of automated systems interacts with the complexity of those systems, thus arguably engendering error discounting in precisely those situations in which it is least advisable. Future work on trust in automation (Lewandowsky, Mundy, & Tan, 2000) is necessary to delineate the circumstances in which such inadvisable error discounting may occur.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

Preparation of this paper was facilitated by a Discovery Grant from the Australian Research Council and an Australian Professorial Fellowship to the second author. The first and third authors were supported by Jean Rogerson Scholarships. Additionally, the third author was supported by an Australian Postgraduate Scholarship and an NIH-NIMH training grant #:T32 MH019879-14. We wish to thank Charles Hanich for assisting with data collection and Gordon Brown for his comments on an earlier version of this manuscript.

References

- Amari S. Theory of adaptive pattern classifiers. *IEEE Transactions in Electronic Computers*. 1967; 16:299–307.
- Baayen RH, Davidson DJ, Bates DM. Mixed-effects mode with crossed random effects for subjects and items. *Journal of Memory and Language*. 2008; 59:390–412.
- Bishop, CM. *Pattern recognition and machine learning*. Cambridge: Springer; 2006.
- Brainard DH. The psychophysics toolbox. *Spatial Vision*. 1997; 10:433–436. [PubMed: 9176952]
- Chater N, Tenenbaum JB, Yuille A. Probabilistic models of cognition: Conceptual foundations. *Trends in Cognitive Science*. 2006; 10:287–291.
- Edgell SE. Delayed exposure to configural information in nonmetric multiple-cue probability learning. *Organizational Behavior and Human Decision Processes*. 1983; 32:55–65.
- Edgell SE, Morrissey JM. Delayed exposure to additional relevant information in nonmetric multiple-cue probability learning. *Organizational Behavior and Human Decision Processes*. 1987; 40:22–38.
- Erickson MA, Kruschke JK. Rules and exemplars in category learning. *Journal of Experimental Psychology: General*. 1998; 127:107–140. [PubMed: 9622910]
- Erickson MA, Kruschke JK. Rule-based extrapolation in perceptual categorization. *Psychonomic Bulletin & Review*. 2002; 9:160–168. [PubMed: 12026949]
- Estes WK. Global and local control of choice behavior by cyclically varying outcome probabilities. *Journal of Experimental Psychology: Learning, Memory, & Cognition*. 1984; 10:258–270.
- Estes WK, Straughn JH. Analysis of a verbal conditioning situation in terms of statistical learning theory. *Journal of Experimental Psychology*. 1954; 47:225–234. [PubMed: 13152299]
- Fantino E, Esfandiari A. Probability matching: encouraging optimal responding in humans. *Canadian Journal of Experimental Psychology*. 2002; 56:58–63. [PubMed: 11901961]
- Farrell S, Lewandowsky S. A connectionist model of complacency and adaptive recovery under automation. *Journal of Experimental Psychology: Learning, Memory, & Cognition*. 2000; 26:395–410.
- Friedman D, Massaro DW. Understanding variability in binary and continuous choice. *Psychonomic Bulletin & Review*. 1998; 5:370–389.
- Friedman, MP.; Burke, CJ.; Cole, M.; Kelly, L.; Millward, RB.; Estes, WK. Two-choice behaviour under extended training with shifting probabilities. In: Atkinson, RC., editor. *Studies in mathematical psychology*. California: Stanford University Press; 1964. p. 250-316.
- Gigerenzer, G.; Todd, PM. *ABC Research Group the. Simple heuristics that make us smart*. New York: Oxford University Press; 1999.
- Gluck MA, Bower GH. Evaluating an adaptive network model of human learning. *Journal of Memory and Language*. 1988a; 27:166–195.
- Gluck MA, Bower GH. From conditioning to category learning: An adaptive network model. *Journal of Experimental Psychology: General*. 1988b; 117:227–24.
- Goldstein, DG.; Gigerenzer, G. the ABC Research Group. The recognition heuristic: how ignorance makes us smart. In: Gigerenzer, G.; Todd, PM., editors. *Simple heuristics that make us smart*. New York: Oxford University Press; 1999. p. 37-58.
- Goldstein DG, Gigerenzer G. Models of ecological rationality: the recognition heuristic. *Psychological Review*. 2002; 109:75–90. [PubMed: 11863042]
- Heskes TM, Kappen B. Learning process in neural networks. *Physical Review A*. 1991; 44:2718–2726. [PubMed: 9906254]
- Histed MH, Pasupathy A, Miller EK. Learning substrates in the primate prefrontal cortex and striatum: sustained activity related to successful actions. *Neuron*. 2009; 63:244–253. [PubMed: 19640482]
- Jones M, Love BC, Maddox WT. Recency effects as a window to generalization: Separating decisional and perceptual sequential effects in category learning. *Journal of Experimental Psychology: Learning, Memory, & Cognition*. 2006; 32:316–332.
- Juslin P, Jones S, Olsson H, Winman A. Cue abstraction and exemplar memory in categorization. *Journal of Experimental Psychology: Learning, Memory, and Cognition*. 2003; 29:924–941.

- Kalish ML, Kruschke JK. Decision boundaries in one-dimensional categorization. *Journal of Experimental Psychology: Learning, Memory, and Cognition*. 1997; 23:1362–1377.
- Kamin, LJ. Predictability, surprise, attention, and conditioning. Campbell, BA.; Church, RM., editors. New York: Appleton-Century-Crofts; 1969. p. 279-296.
- Kerstholt JH, Passenier PO. Fault management in supervisory control: the effect of false alarms and support. *Ergonomics*. 2000; 43:1371–1389. [PubMed: 11014759]
- Kruschke JK. ALCOVE: An exemplar-based connectionist model of category learning. *Psychological Review*. 1992; 99:22–44. [PubMed: 1546117]
- Kruschke JK, Blair NJ. Blocking and backward blocking involve learned inattention. *Psychonomic Bulletin & Review*. 2000; 7:636–645. [PubMed: 11206204]
- Kruschke JK, Johansen MK. A model of probabilistic category learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*. 1999; 25:1083–1119.
- Lewandowsky S, Mundy M, Tan G. The dynamics of trust: Comparing humans to automation. *Journal of Experimental Psychology: Applied*. 2000; 6:104–123. [PubMed: 10937315]
- Little DR, Lewandowsky S. Better learning with more error: probabilistic feedback increases sensitivity to correlated cues in categorization. *Journal of Experimental Psychology: Learning, Memory, & Cognition*. 2009a; 35:1041–1061.
- Little DR, Lewandowsky S. Beyond non-utilization: irrelevant cues can gate learning in probabilistic categorization. *Journal of Experimental Psychology: Human Perception & Performance*. 2009b; 35:530–550. [PubMed: 19331506]
- Lorenz, B.; Parasuraman, R. Automated and interactive real-time systems. In: Durso, FT.; Nickerson, R.; Dumais, S.; Lewandowsky, S.; Perfect, T., editors. *Handbook of applied cognition*. Chichester, England: Wiley; 2007. p. 415-441.
- McKinley SC, Nosofsky RM. Investigations of exemplar and decision bound models in large, ill-defined category structures. *Journal of Experimental Psychology: Human Perception and Performance*. 1995; 21:128–148. [PubMed: 7707028]
- Moray, N. Industrial systems. In: Durso, FT.; Nickerson, R.; Dumais, S.; Lewandowsky, S.; Perfect, T., editors. *Handbook of applied cognition*. Chichester, England: Wiley; 2007. p. 281-305.
- Murata N, Kawanabe M, Ziehe A, Müller K-R, Amari S. On-line learning in changing environments with applications to supervised and unsupervised learning. *Neural Networks*. 2002; 15:743–760. [PubMed: 12371524]
- Myung I-J, Pitt MA. Model comparison methods. *Methods in Enzymology*. 2004; 383:351–366. [PubMed: 15063657]
- Nosofsky RM. Attention, similarity and the identification-categorization relationship. *Journal of Experimental Psychology: General*. 1986; 115:39–57. [PubMed: 2937873]
- Parasuraman R, Riley V. Humans and automation: use, misuse, disuse, abuse. *Human Factors*. 1997; 39:230–253.
- Pelli DG. The videotoolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision*. 1991; 10:437–442. [PubMed: 9176953]
- Rescorla, RA.; Wagner, AR. A theory of pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In: Black, AH.; Prokasky, WF., editors. *Classical conditioning: II. current research and theory*. New York: Appleton-Century-Crofts; 1972. p. 64-99.
- Rouder JN, Ratcliff R. Comparing categorization models. *Journal of Experimental Psychology: General*. 2004; 133:63–82. [PubMed: 14979752]
- Rumelhart, DE.; Hinton, GB.; Williams, RJ. Learning internal representations by back-propagation errors. In: Rumelhart, DE.; McClelland, JL., editors. *Parallel distributed processing: Explorations in the microstructure of cognition*. Vol. 1. Cambridge, MA: MIT Press; 1986. p. 318-362.
- Shanks DR, Tunney RJ, McCarthy JD. A re-examination of probability matching and rational choice. *Journal of Behavioral Decision Making*. 2002; 15:233–250.
- Smith TA, Kimball DR. Learning from feedback: spacing and delay-retention effects. *Journal of Experimental Psychology: Learning, Memory, & Cognition*. 2010; 36:80–95.
- Stewart N, Brown GDA. Sequence effects in the categorization of tones varying in frequency. *Journal of Experimental Psychology: Learning, Memory, and Cognition*. 2004; 30:416–430.

- Stewart N, Brown GDA, Chater N. Sequence effects in categorization of simple perceptual stimuli. *Journal of Experimental Psychology: Learning, Memory, and Cognition*. 2002; 28:3–11.
- Tversky A, Kahneman D. Judgment under uncertainty: Heuristics and biases. *Science*. 1974; 185:1124–1131. [PubMed: 17835457]
- Wickens, C. Aviation. In: Durso, FT.; Nickerson, R.; Dumais, S.; Lewandowsky, S.; Perfect, T., editors. *Handbook of applied cognition*. Chichester, England: Wiley; 2007. p. 361-389.
- Yelen DR, Yelen D. Overtaining reversal effect in a probability learning situation. *Psychonomic Science*. 1969; 14:152.

Figure 1.

Probability of responding with category A for each of the training items across training blocks for early, mid, and late conditions in Experiment 1. The legend identifies each item's reinforcement probabilities, $P(A|j)$, before and after the shift.

Figure 2. Observed slope through response probabilities across all 4 training items in Experiment 1 (solid lines) and objective slopes (dotted lines). Error bars indicate 95% confidence intervals. The three panels show the early (top), mid (middle), and late (bottom) condition, respectively.

Figure 3.

Decisional-recency analysis for the mid and late conditions in Experiment 1. The probability of shifting the response from one category to the other on trial n is shown as a function of whether the response on trial $n - 1$ was an error or correct. The top panel is for two identical stimuli in succession and the bottom panel is for two maximally dissimilar stimuli following each other. The smooth lines are based on locally-weighted regression (*lowess*). See text for details of alignment of conditions.

Figure 4.

Observed slope through response probabilities across all 4 training items in Experiment 2 (solid lines) and objective slopes (dotted lines). Error bars indicate 95% confidence intervals. The two panels show the early (top) and late (bottom) condition, respectively.

Figure 5.

Decisional-recency analysis for all conditions in Experiment 2. The probability of shifting the response from one category to the other on trial n is shown as a function of whether the response on trial $n - 1$ was an error or correct. The top panel is for two identical stimuli in succession and the bottom panel is for two maximally dissimilar stimuli following each other. The smooth lines are based on locally-weighted regression (*lowess*). See text for details of alignment of conditions.

Table 1

Results of multilevel generalized linear model analysis of trial-by-trial contingencies from the mid and late conditions of Experiment 1 and Experiment 2.

Fixed Effect ^a	Estimate	z-value	p
perceptual difference (PD)	4.21	15.31	< .0001
error before (EB)	2.53	12.60	< .0001
block ^b	0.01	0.58	n.s.
PD × EB	-3.92	-14.83	< .0001
PD × block	.02	0.79	n.s.
EB × block ^c	-.07	-3.33	< .001
PD × EB × block	.07	1.86	n.s.

^aThe best-fitting model ($BIC = 6693$) also included subject, EB, PD, and block as random effects.

^blinear effect of block when the preceding response was correct ($EB = 0$)

^cadjustment to linear effect of block when the preceding response was incorrect ($EB = 1$)

Negative log-likelihood ($-\ln L$), corrected Akaike Information Criterion values, (AIC_c), AIC weights, $w(AIC)$, Bayesian Information Criterion values, (BIC), and BIC weights, $w(BIC)$, for fits of the GCM, the MAC model, RASHNL and *r*ATRIUM across both experiments and all participants.

Table 2

Model	Annealing	Free params	$-\ln L$	AIC_c	$w(AIC)$	BIC	$w(BIC)$
GCM	-	2	27295.92	54274.55	0	55147.81	0
MAC	On	2	27601.90	54880.43	0	55751.22	0
MAC	Off	1	28875.47	59067.51	0	58024.65	0
RASHNL	On	4	25256.81	50315.53	1	51625.56	1
RASHNL	Off	3	25774.22	51321.99	0	52382.39	0
<i>r</i> ATRIUM	On	4	25637.00	51075.90	0	52385.92	0
<i>r</i> ATRIUM	Off	3	26293.77	52361.10	0	53421.50	0

Note: Bold items indicate best fit.