UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

More than One Kind of Probability Matching: Evidence from a Dual-Task Paradigm

Permalink

Journal

ISSN

Authors

Otto, A. Ross
Markman, Arthur
Taylor, Eric

Publication Date

2010

Peer reviewed

More than One Kind of Probability Matching: Evidence from a Dual-Task Paradigm

A. Ross Otto and Arthur B. Markman
(rotto@mail.utexas.edu, markman@psy.utexas.edu)
Department of Psychology, University of Texas, Austin, TX 78712 USA

Eric G. Taylor (etaylor4@illinois.edu)
Department of Psychology, University of Illinois, Champaign, IL 61820 USA

Abstract

Probability-matching is a well-documented suboptimal behavior that arises in simple prediction tasks. We identify two distinct, local choice strategies that both give rise to probability-matching behavior on a global level. Using a dual-task paradigm, we evaluate the hypothesis that these qualitatively different strategies exhibit different demands on individuals' central executive resources. We find that participants placed under a concurrent working memory are driven away from the one-trial-back strategy—utilized by participants without a working memory load—and towards a strategy that integrates a longer window of past outcomes into the current prediction. In other words, the demands of the concurrent task appeared to shift the prediction strategies used by decision-makers in our study.

Keywords: Decision-making; Prediction; Win-Stay-Lose-Shift; Working Memory; Dual Task; Heuristics

Introduction

One decision-making anomaly of great interest is the tendency for humans to match their responses to outcome probabilities in the prediction of binary outcomes. For example consider a laboratory task in which people need to repeatedly predict which of two outcomes (say Event A and Event B) will occur next. If Event A occurs at a base rate of $p = .65$, Event B occurs at a base rate of $p = .35$ and each outcome is conditionally independent of the last outcome, the optimal prediction strategy would be to always predict that Event A will occur next, which is called *maximizing*. However, a large body of empirical work suggests that people appear to predict events in proportion to their frequency of occurrence, known as *probability matching* (Estes, 1961; Vulkan, 2000). Under probability matching, a person would predict Event A 65% of the time and Event B 35% of the time. It is easy to see that this strategy produces an expected overall accuracy of 54.5% (calculated as .65 x .65 + .35 x .35), which is inferior to that produced by *maximizing*—which produces an expected overall prediction accuracy of 65%. In the present study, we examine strategies that be may underlying probability matching in random sequences of events.

The psychological mechanisms that give rise to probability matching behavior are unclear and are a matter of ongoing debate. One hypothesis posits that probability matching arises from the use of a suboptimal cognitive shortcut in which individuals allocates their responses according to an assessment of the observed outcome probabilities (e.g., Koehler & James, 2009). Under this strategy, termed *expectation matching* (EM), the decision-maker's responses are the

result of integrating a moving window of past outcome information (Sugrue, Corrado, & Newsome, 2004). To generate a response, the individual stochastically and independently generates predictions in accordance with this historical assessment of outcome probabilities. Assuming a sufficiently long historical window, a decision-maker utilizing the EM strategy in the example above would stochastically allocate 65% of their predictions to Event A and 35% of their predictions to Event B.

Another proposal suggests that probability matching behavior seen at a more global level is the byproduct of a local decision process called *win-stay lose-shift* (WSLS: Herrnstein, Rachlin, & Laibson, 2000). Under WSLS, an individual persists with predicting one event, say Event A, until they make an incorrect prediction, at which point they shift responses and persist with predicting Event B until they are incorrect. While under certain task circumstances WSLS is an optimal choice strategy (Shimp, 1976), it is a suboptimal prediction strategy in the task outlined above. It can be shown that WSLS produces overall response rates (and hence, accuracy rates) equivalent to probability matching (Unturbe & Corominas, 2007). Further, there is evidence that people utilize WSLS in the simple binary prediction task described above (Gaissmaier & Schooler, 2008). Unlike the EM strategy, which involves integrating a comparatively long historical window of outcomes, WSLS requires that the decision-maker maintain a short-term memory for only the most recent response and outcome.

In the present study, we examined the cognitive demands imposed by the WSLS and EM strategies, with the idea that decision makers may utilize both strategies, but under different circumstances. While both strategies result in equivalent behavior at a global level—probability matching—they make different behavioral predictions at a local, trial-by-trial level. It is well documented that the working memory demands of a secondary task deplete mental resources that could otherwise be used to accomplish a primary task (Pashler, 1994). For example, Zeithamova and Maddox (2006) found that working memory load disrupts learning of explicit, rule-based categories and drives participants towards the use of an implicit, information-integration strategy. Here, we place decision-makers under a concurrent working memory load and find that they exhibit the same global tendency to probability match as decision-makers without a working memory

load. Using simple models, we demonstrate that different local strategies result in global probability matching. The distinction between these two matching strategies is theoretically significant because recent contributions to the probability matching literature (e.g., Gaissmaier & Schooler, 2008; Koehler & James, 2009) fail to find common ground on a) which strategies may give rise to probability matching behavior, and b) to what extent these strategies place demands on executive function.

Method

Participants   One-hundred and sixty undergraduates at the University of Texas at Austin participated in this study, randomly assigned to one of two conditions: Dual-Task (DT) and Single-Task (ST). Participants were paid a small cash bonus of one cent per correct prediction.

Design and Procedure   The experiment stimuli and instructions were displayed on 17-inch monitors. The participants were told that their goal was to predict repeatedly whether a red square would appear above a fixation cross or a green square below the fixation cross, using the up and down arrows respectively (see Figure 1 for a task screenshot). Like other studies (e.g., Koehler & James, 2009), the sequence of events was serially independent. The probability of the more common event was $p = .65$. The assignment of the high-probability event to the outcomes was counterbalanced across subjects. Subjects completed 10 practice trials in order to familiarize themselves with the response procedure, followed by 320 trials divided into 8 blocks of 40 trials each.


Figure 1: Example task screenshot of response and outcome for a correct prediction.

In order to accommodate the dual-task manipulation, the prediction task used a deadline procedure to ensure that a fixed amount of time elapsed each trial. At the start of each trial, the subject saw the word "PREDICT" and had two seconds to make a response. This response window lasted two seconds regardless of the timing of the response, and was followed by the actual outcome along with feedback indicating whether their prediction was correct ("CORRECT") or incorrect ("INCORRECT"). The outcome and feedback were displayed for one second, and was followed by a one second inter-trial interval. If a subject failed to respond within the response window, the message "TOO SLOW" was displayed

along with the outcome. The timing of response windows and outcomes was the same for both the ST and DT conditions.

Blocks in the DT condition consisted of a secondary tone-counting task in addition to the prediction task. The design of the secondary task follows that of Foerde et al. (2007). Two types of tones, high-pitched (1000 Hz) and low-pitched (500 Hz) were played during each trial in the DT condition. Each three-second trial was divided into 12 intervals of 250 ms, with the tones occurring in intervals 3-10 (500-2,500 ms after trial onset). The number of tones presented each trial varied uniformly between 1 and 3 and occurred randomly within intervals 3-10. The pitch of each tone varied randomly, with the base rate of high tones varying uniformly from .3 to .7 each block. The subjects were instructed to maintain a running count of the number of high tones while ignoring the low-pitched tones. Note that the secondary task persisted during both the response window and the outcome. At the end of each 40-trial block, the subjects reported their running count using the keyboard and were instructed to restart their count at zero.

After subjects had completed 320 trials, they completed a questionnaire in which they were asked to provide estimates of the overall frequency of the red and green events. They were also given five prediction strategies to evaluate. These strategies included an expectation matching strategy ("Predict GREEN 65% of the time regardless of what happened during the last outcome"), a maximizing strategy, ("Always predict GREEN, regardless of what happened during the last outcome"), and a WSLS strategy ("Stick with predicting one outcome, and then change your prediction if you were incorrect on the last trial"). Subjects were instructed to rank these five strategies from 1 ("the best possible strategy") to 5 ("the worst possible strategy"), using each ranking only once.

Results

We removed data from 12 ST and 26 DT participants whose prediction behavior differed non-significantly from equiprobable responding (Binomial test at the $p = .05$ level of significance). We also removed the data of eleven participants who failed to respond before deadline more than 20 times during the experiment. One hundred and eleven participants (48 DT and 63 ST participants) remained in the analysis that follows.

Overall Prediction Performance   Figure 2 depicts the subjects' accuracy, by condition, in predicting outcomes over the 320 trials. The dashed line depicts the level of accuracy expected under probability matching probability—that is, if participants allocated their 65% of their responses to the more frequent outcome. A 2 (task condition) x 2 (trial block) ANOVA revealed neither a significant main effect of task condition, $F(1,107) = .55$, $p = .46$, nor a significant interaction between condition and trial block, $F(1,107) = 0.27$, $p = .61$. There was a significant main effect of trial block, $F(1,107) = 25.51$, $p < .001$. Again, the lack of effect of task condition suggests that the dual task manipulation did not hinder subjects' overall accuracy, but rather, may have shifted the

Figure 2: Left panel: mean prediction accuracy, by task condition and trial block. ST = Single-task condition, DT=dual-task condition. Error bars represent standard error of the mean.


Figure 3: Proportion of Subjects Deviating Significantly from Matching (by Binomial test), by task condition and trial block. ST = Single-task condition, DT=dual-task condition. Error bars represent standard error of proportion.

prediction strategies that subjects employed.

Overall Deviation from Matching Recall that our main goal was to determine whether matching behavior results from different strategies across the ST and DT conditions. Before comparing strategy usage, we first determine that both groups were in fact predominantly matching—and to the same degree. Specifically, we determined whether the secondary task manipulation affected the degree to which subjects deviated significantly from matching behavior (that is, allocating 65% of one's responses to the more frequent event). For each of the 8 blocks, we calculated the proportion of subjects whose response allocations deviated significantly from a response allocation that matched the observed outcome frequency. The proportion of subjects in each condition, by block, that deviated significantly from probability matching behavior (under a Binomial test at the $p = .05$ level significance) are shown in Figure 3. We conducted a logistic regression with each subject's classification (deviating significantly or not) as the criterion and task condition and trial block as predictors, observing no significant coefficients for task condition (*Beta* = -.83, $p = .44$) or the interaction between task condition and trial block (*Beta* = .08, $p = .53$). Trial block did have a significant coefficient (*Beta* = .5, $p < .001$). The apparent null effect of task condition suggests that ST and DT subjects were engaging in prediction behavior that appears similar at a coarse level of analysis.

Exponentially-Weighted Averaging Model Analysis At least two distinct response strategies can manifest themselves as probability matching. Under WSLS, the decision-maker repeats the previous trial's response after a correct prediction and switches their response after an incorrect prediction. Thus responses under WSLS are determined by the outcome

on the only the most recent trial. In contrast, EM requires that the decision-maker integrate a much longer window of previous outcomes, which in turn informs the decision-maker's response probabilities. By fitting a simple exponentially-weighted averaging model model to participants' responses, we identified the degree to which participants' predictions were dependent on recent outcomes. The probability P(t) of the decision-maker predicting the green event at time t is determined by:

$$P(t) = recency*outcome(t\text{-}1) + (1\text{-}recency)*P(t\text{-}1),$$

where *outcome(t-1)* is the outcome on the previous trial, *P(t-1)* is the model's estimate of the rate at which the green outcome occurs, and *recency* is a parameter that determines how much recent outcomes are weighted in updating *P(t)*. When the recency parameter is large, *P(t)* is based only on the most recent trial's outcome, and when the recency parameter is small, the model's predicated response on the next trial *P(t)* is based on a long window of previous outcomes. We fit this model to each participants' responses using maximum likelihood estimation, assuming separate parameter values across blocks. As shown in Figure 4, ST participants had larger estimated learning weights than DT participants, indicating that prediction strategies employed by ST participants were influenced more by recent outcomes. A 2 (task condition)x2 (trial block) ANOVA revealed a significant main effect of task condition, $F(1,107) = 4.13$, $p < .05$, a significant main effect of block, $F(1,107) = 21.38$, $p < 0.001$, and a significant interaction between condition and trial block, $F(1,107) = 6.34$, $p < .05$. The effect of condition suggests that ST participants exhibited choice behavior characteristic of WSLS—dependence on only the most recent trials—while DT participants used a strategy characteristic of the EM strategy—involving integration of a long window of past outcomes.

Figure 4: Average best-fitting recency parameter values for exponentially-weighted averaging model, by task condition and block. ST = Single-task condition, DT=dual-task condition. Error bars represent standard error of the mean.
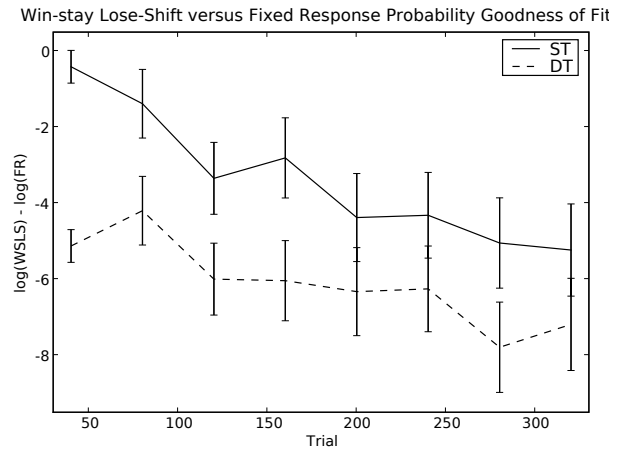

Figure 5: Comparison of model goodness-of-fit between WSLS and EM models. Average likelihood ratios using best-fitting parameter values for each block of each subject. Error bars represent standard error of the mean. ST = Single-task condition, DT=dual-task condition. Error bars represent standard error of the mean.

Models of the Two Prediction Strategies To more directly address usage of these strategies, we compared the relative goodness-of-fit of two models that instantiated the WSLS and EM strategies. To examine participants' WSLS usage, we fit a simple WSLS model to participants' choices, hypothesizing that ST participants would be better fit by this model than DT participants. This one-parameter model constrains the probability of a switching responses after an incorrect prediction (or a "loss") to the probability of persisting with the same response after a correct response (or a "win"). This model follows the WSLS implementation described by Steyvers, Lee, and Wagenmakers (2009). To examine usage of the EM strategy, we fit a simple stochastic response model, which we call the fixed response probability (FR) model, to participants' data. Under this model, a single parameter determines the base rate of predicting the green event. This model—which we use a proxy measure for EM strategy use—assumes that responses are determined stochastically and independently. One crucial difference between these two models is the dependence of the response on trial t to the outcome on trial t-1. We fit both models to each participants' choice data using maximum likelihood estimation allowing parameter values to vary across blocks.

We predicted that ST subjects would be better described by the WSLS model and that DT subjects would be better described by the FR model. Figure 5 depicts the relative goodness-of-fit (expressed as a log-likelihood ratio) between the two models, for each condition across the 8 blocks. Indeed, the likelihood ratios reveal that ST participants were better described by the WSLS model than the responses of DT participants, and conversely, DT participants were better described by the FR model—our proxy for the EM strategy. A 2 (task condition) x 2 (trial block) ANOVA revealed a significant main effect of task condition, $F(1,107) = 5.28, p<.05,$

a main effect of trial block, $F(1,107) = 19.18, p<.001,$ and no significant interaction between task condition and trial block, $F(1,107) = 1.14, p=.29.$ The main effect of task condition suggests that the concurrent working memory load influenced the local prediction strategies utilized by decision-makers.

Offline Reported Event Probabilities We hypothesized that the secondary task would impair DT participants' ability to explicitly encode information about outcome frequencies. To test this, we calculated absolute deviations between participants' offline reported outcome probabilities and true empirical base rates. The average absolute deviations are shown in Figure 6. We found that DT participants' reported outcome probabilities deviated significantly more from observed base rates than DT participants, $t(107) = 2.82, p<.01.$ Taken in conjunction with the similar overall accuracy profiles of the two groups, this result suggests that the two groups may have been using qualitatively different strategies to make predictions.

Strategy Self-Reports We assessed participants' offline endorsement of the strategies that were described in the questionnaire. To do this, we compared participants' relative preference for the WSLS over EM by their subtracting their ranking of the WSLS strategy from their ranking of the EM strategy, yielding a measure of endorsement of WSLS (note that this measure is equally informative about preference for EM). We found that ST participants' endorsement of WSLS significantly correlated with their overall WSLS model goodness-of-fit, $r(107) = .35, p < .01,$ suggesting that ST participants had some explicit awareness of the strategies they employed. In contrast, DT participants' strategy endorsements did not significantly correlate with their average goodness-of-fit mea-

Figure 6: Mean absolute deviation from observed (empirical) base rate, by task condition. Error bars represent standard error of the mean.

sures for either model, suggesting the concurrent working memory load impaired decision-makers' ability to explicitly report the strategies they employed [WSLS model: $r(107) = .15$, $p=.28$, FR model: $r(107) = -.02$, $p = .82$].

Discussion

In this experiment, we investigated the effect of a concurrent working memory task on probability-matching behavior in a random, sequentially independent prediction task. To do so, we imposed a secondary working memory task on subjects, which was believed to deplete working memory resources that could have been used on the primary prediction task (Pashler, 1994). In the DT condition, subjects needed to both make responses in the prediction task and update their count of auditory tones, while in the ST condition, subjects needed only to make predictions. Although most subjects in both conditions demonstrated probability matching, subjects in the ST condition relied more on a WSLS strategy, which requires memory for the previous prediction and outcome. This finding suggests that while both ST and DT subjects appear to be using suboptimal strategies with similar base rates at a molar level, the two groups may actually be using different prediction strategies.

Our results are interesting in the context of previous dual-task studies of human learning. For example, Foerde et al., (2007) found that a concurrent working memory load during probabilistic classification learning impaired subjects' acquisition of explicit associations between perceptual cues and outcomes, although these subjects evidenced implicit learning of cue-outcome contingencies. Further, they were unable to flexibly apply knowledge about cues in an offline evaluation. Zeithamova and Maddox (2006) found that a concurrent working memory load disrupts learning of explicit, rule-based categories and instead drives subjects towards the use of an implicit, information integration strategy. Both of these studies point to the possibility that concurrent working memory load engenders the use of implicit learning systems. In our study, utilization of the EM strategy may be indicative of the operation of an implicit system.

Another possibility raised in the literature is that probabil-

ity matching arises out of peoples' search for regularities in the event sequences (Gaissmaier & Schooler, 2008). Even when laboratory prediction tasks are probabilistic and outcomes sequences are conditionally independent, people may search for deterministic patterns in an attempt to achieve prediction accuracy above that of maximizing. Thus, if an individual believes that the event sequence contains structure, he or she will try to improve their accuracy by searching for patterns. Gaissmaier & Schooler's result suggests that that some individuals in the present study who appear to be probability matching—rather than maximizing—are more adept at detecting deterministic patterns when they are later introduced into the sequence of events.

One possibility in the present study is that subjects in the ST condition may have begun a search for deterministic patterns and abandoned the search given the very low likelihood of a pattern repeating itself in the random sequence, reverting later to a suboptimal WSLS strategy. Supporting evidence comes from the fact that over 60% of the ST condition's responses are consistent with WSLS and that this percentage increases over time. This hypothesis will be the subject of investigation in future studies.

References

Foerde, K., Poldrack, R. A., & Knowlton, B. J. (2007). Secondary-task effects on classification learning. *Memory & Cognition*, *35*(5), 864–874

Gaissmaier, W., & Schooler, L. J. (2008). The smart potential behind probability matching. *Cognition*, *109*(3), 416–422

Herrnstein, R. J., Rachlin, H., & Laibson, D. I. (2000). *The matching law*. Harvard University Press.

Koehler, D. J., & James, G. (2009). Probability matching in choice under uncertainty: Intuition versus deliberation. *Cognition*, *113*(1), 123–127

Pashler, H. (1994). Dual-task interference in simple tasks: Data and theory. *Psychological Bulletin*, *116*(2), 220–244

Shimp, C. P. (1976). Short-term memory in the pigeon: the previously reinforced response. *Journal of the Experimental Analysis of Behavior*, *26*(3), 487–493

Steyvers, M., Lee, M. D., & Wagenmakers, E. (2009). A bayesian analysis of human decision-making on bandit problems. *Journal of Mathematical Psychology*, *53*(3), 168–179

Sugrue, L. P., Corrado, G. S., & Newsome, W. T. (2004). Matching behavior and the representation of value in the parietal cortex. *Science*, *304*(5678), 1782–1787

Unturbe, J., & Corominas, J. (2007). Probability matching involves rule-generating ability: A neuropsychological mechanism dealing with probabilities. *Neuropsychology*, *21*(5), 621–630

Vulkan, N. (2000). An economist's perspective on probability matching. *Journal of Economic Surveys*, *14*(1), 101–118

Zeithamova, D., & Maddox, W. T. (2006). Dual-task interference in perceptual category learning. *Memory & Cognition*, *34*(2), 387–398