There are at Least Two Kinds of Probability Matching: Evidence from a Secondary Task

A. Ross Otto
Department of Psychology
University of Texas at Austin

Eric G. Taylor
Department of Psychology
Yale University

Arthur B. Markman
Department of Psychology
University of Texas at Austin

Cognition (in press)

Please address all correspondence to:
A. Ross Otto
Department of Psychology
University of Texas
1 University Station, A8000
Austin, TX 78712
Phone: (512) 232-2484
Fax: (512) 471-6175
rotto@mail.utexas.edu

Abstract

Probability matching is a suboptimal behavior that often plagues human decision-making in simple repeated choice tasks. Despite decades of research, recent studies cannot find agreement on what choice strategies lead to probability matching. We propose a solution, showing that two distinct local choice strategies—which make different demands on executive resources—both result in probability-matching behavior on a global level. By placing participants in a simple binary prediction task under dual- versus single-task conditions, we find that individuals with compromised executive resources are driven away from a one-trial-back strategy (utilized by participants with intact executive resources) and towards a strategy that integrates a longer window of past outcomes into the current prediction. Crucially, both groups of participants exhibited probability-matching behavior to the same extent at a global level of analysis. We suggest that these two forms of probability matching are byproducts of the operation of explicit versus implicit systems.

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, imagine a series of horse races where one needs to predict which of two Horses will win. If Horse A wins 65% of the time and Horse B wins 35% of the time, and each race is conditionally independent of the last, the optimal prediction strategy would be to predict that Horse A will win every race. This prediction strategy is called *maximizing*.

A large body of research suggests that people predict events in proportion to their frequency of occurrence, a strategy known as *probability-matching*. Over the years, the cognitive mechanisms underlying probability-matching have been the topic of much curiosity and speculation (for a review, see Vulkan, 2000). Under probability-matching, a person predicts Horse A 65% of the time and Horse B 35% of the time. It is easy to see that this strategy produces an expected overall accuracy of 54.5% ($.65\times.65+.35\times.35$), which is inferior to that produced by maximizing—which produces an expected accuracy of 65%. In the present study, we examine strategies that may underlie probability-matching in independent event sequences.

The psychological mechanisms that give rise to probability-matching behavior are unclear and are a matter of ongoing debate. One hypothesis is that probability-matching arises from the use of a heuristic under which individuals allocate their responses according to an assessment of the observed outcome probabilities (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 horse-racing example would stochastically allocate 65% of their predictions to Horse A and 35% of their predictions to Horse B.

Another proposal is 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, individuals persist in 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. Under certain task circumstances WSLS is an optimal choice strategy but it is a suboptimal prediction strategy in the task outlined above, producing overall response rates—and hence, accuracy rates—equivalent to probability-matching (Shimp, 1976). 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 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. Indeed, the proposal that different local choice behaviors can give rise to global matching behavior has been considered in the animal literature (Hinson & Staddon, 1983).

It is well documented that the working memory (WM) demands of secondary tasks deplete mental resources that could be used to accomplish primary tasks (Pashler, 1994). For example, Filoteo, Lauritzen, and Maddox (2010) found that WM 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 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 (Gaissmaier & Schooler, 2008; Koehler & James, 2009).

One possibility is that WSLS requires maintaining an active representation of the outcome and response from the most recent trial and thus imposes a burden on WM and executive resources, while EM entails gradual integration of outcome information, and as such, may be characteristic of an implicit learning mechanism (Ashby et al., 1998). Because research suggests that WM load drives individuals towards the use of implicit or procedural strategies in classification (Foerde, Knowlton, & Poldrack, 2006; Filoteo et al., 2010), and these strategies characteristically entail accrual of information over many trials, we predict that WM load should drive participants towards a strategy like EM. Accordingly, we expect participants without working memory load to exhibit comparatively greater reliance on WSLS, which characteristically involves a shorter window of outcome integration.

We provide converging evidence that WM load engenders usage of implicit strategies by examining participants' explicit encoding of environmental outcome probabilities and their self-reported strategy usage. Category learning work suggests that use of rule-based strategies—which are thought to rely on explicit processes—can be accurately verbalized in participants' self-reports, while information-integration strategies—which are thought to rely on implicit processes—are not accurately described by participants (Ashby & Maddox, 2005). As such, we do not expect participants' self-reports, regardless of WM load, to predict model-assessed EM strategy use. We expect that participants without WM load will accurately report WSLS use because they should be able to form declarative knowledge of strategy use.

Method

Participants

160 students participated for course credit, assigned to one of two conditions: Dual-Task (DT) and Single-Task (ST), and were paid one cent per correct prediction.

Design and Procedure

The main task was predict whether a red square would appear above a fixation cross or a green square would appear below the fixation cross, using the keyboard. The sequence of events was serially independent, and the probability of the more common event was p=.65. Participants completed 10 practice trials to familiarize themselves with the timed procedure, followed by 320 trials.

The prediction task used a deadline procedure to ensure that a fixed amount of time elapsed each trial. On each trial, participants saw the word "PREDICT" and had two seconds to make a response, after which the actual outcome and visual feedback ("CORRECT" or "INCORRECT") was displayed for one second, followed by a one-second inter-trial interval. The timing of trials was equivalent between ST and DT conditions.

In the DT condition, two types of tones, high-pitched (1000 Hz) and low-pitched (500 Hz) were played during each trial. Each three-second trial was divided into 12 intervals of 250 ms, with tones occurring in intervals 3-10 (500-2,500ms after trial onset). The number of tones presented each trial varied uniformly between 1 and 3 occurring randomly within intervals 3-10. The base rate of high tones was determined every 40 trials, varying uniformly between .3 and .7. Participants were instructed to maintain a running count of the number of high tones while ignoring the low-pitched tones. At the end of each 40-trial block, participants reported their count and were instructed to restart their count at zero.

After the main task, participants completed a questionnaire in which they were asked to indicate the probability of red and green events. They were also given five prediction strategies to evaluate in a hypothetical environment with a given event base rate (67%). These strategies included an EM strategy ("Predict GREEN 67% 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"). Participants were instructed to rank these five strategies from one (best) to five (worst).

Results

To ensure that we analyzed the behavior of participants who exhibited sensitivity to outcomes, we removed data from 11 ST and 26 DT participants who allocated less than 50% of their responses to more frequent response. One hundred and twenty-three participants (53 DT, 70 ST) remained in the analysis that follows.

Reported Event Base Rates

We calculated absolute deviations between participants' offline reported outcome probabilities and true empirical base rates, finding that DT participants' reported outcome probabilities (M=.10,SD=.07) deviated significantly more from observed base rates than ST participants (M=.07,SD=.05)[$t(121)=2.21$,$p<.05$,d=.44]. The apparent difference in explicit encoding of outcome frequencies suggests that the secondary task impaired DT participants' ability to explicitly encode information about outcome frequencies.

Figure 1.

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. Right panel: proportion of participants deviating significantly from probability-matching (by Binomial test at p=.05 level of significance) by task condition and trial block. Error bars represent standard error of proportion.

A B

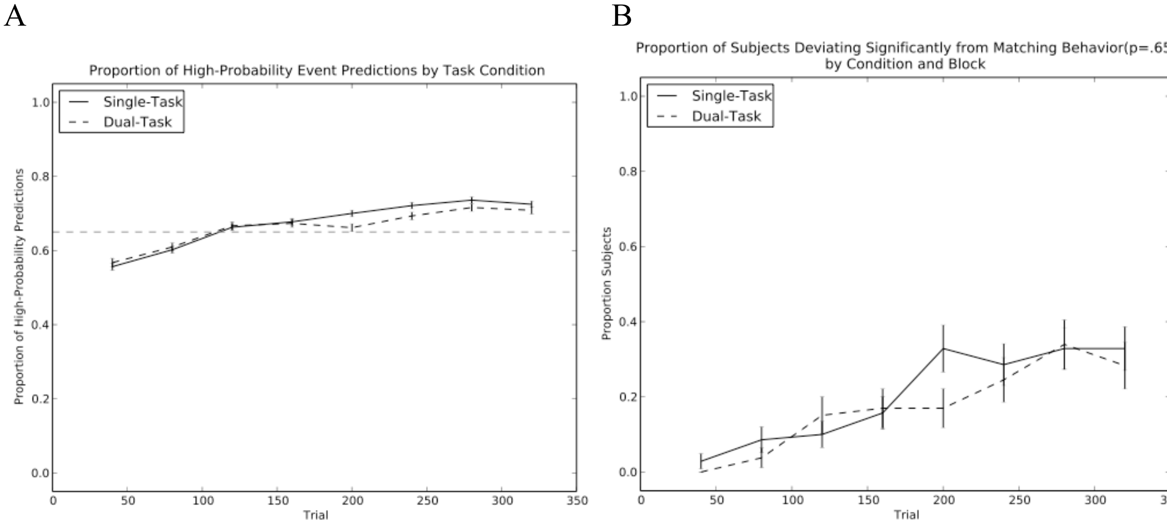

Overall Response Rates

Figure 1A depicts the participants' overall rates of predicting the more common event over 320 trials. The dashed line depicts probability-matching—that is, if participants allocated their 65% of their responses to the more frequent outcome. A 2 (condition) x 2 (trial block) ANOVA revealed neither a significant main effect of condition, $F(1,121)=.26, p=.61, \eta_p^2=.002$, nor a significant interaction between condition and trial block, $F(1,121)=1.35, p=.25, \eta_p^2=.011$, revealing that the dual task manipulation did not significantly alter participants' tendency to probability-match. There was a significant main effect of trial block, $F(1,121)=80.71, p<.001$, $\eta_p^2=.667$, suggesting that both groups may have begun to exhibit response rates greater than probability-matching, mirroring previous work (Fantino & Esfandiari, 2002).

Deviation from Probability Matching

Our main goal is to determine whether matching behavior arises 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 examined whether the secondary task manipulation affected the degree to which participants deviated significantly from matching behavior (that is, allocating 65% of one's responses to the more frequent event). For each of the eight blocks, we calculated the proportion of participants whose response allocations deviated significantly from 65%. The proportion of participants in each condition, by block, that deviated significantly from probability matching behavior (Binomial test,$\alpha=.05$[1]) are shown in Figure 1B. We conducted a hierarchical logistic regression predicting the log odds of deviating with condition and block as predictors. The effect of block was significant (the odds of deviating were between 1.65 and 2.03 times greater each successive block,$\alpha=.05$), but neither condition ($\gamma=-.56, SE=.92, p=.57$), nor the interaction between condition and trial block ($\gamma=-.28, SE=.12, p=.78$) were significant predictors. The apparent null effect of task condition suggests that the level of matching behavior did not differ between ST and DT participants.

Recency-Weighted Average Model

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 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 recency-weighted learning model (Rescorla& Wagner, 1972) to participants' choices we identified the degree to which predictions depended on recent outcomes. The probability P of the decision-maker predicting the green event on trial t+1 is determined by

$$P_{t+1} = P_t + \alpha\left(O_t - P_t\right) \tag{1}$$

where O_t is the outcome on the previous trial, P_t is the previous estimate of the rate at which the green outcome occurs, and α is a parameter that determines how much recent outcomes are weighted in updating P. When α is large, P_{t+1} is based only the outcome on trial t, and when α is

[1] Under this threshold, deviating implies 31 or more predictions of the more frequent event (of 40).

small P_{t+1} is based on a long window of previous outcomes[2]. We fit this model to each participant's responses using maximum likelihood, assuming separate parameter values across blocks. As shown in Figure 2, 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 (condition)x8 (block) ANOVA revealed a significant main effect of task condition, $F(1,121)=5.28, p<.05, \eta_p^2=.044$, but no effect of block, $F(1,121)=2.30, p=.13, \eta_p^2=.019$ and a no interaction between condition and trial block, $F(1,121)=2.50, p=.12, \eta_p^2=.021$, suggesting ST participants exhibited choice behavior more characteristic of WSLS—dependence on only the most recent trials—while DT participants used a strategy characteristic of the EM strategy—involving integration a longer window of past outcomes.

Figure 2.

Average best-fitting recency parameter values for recency-weighted average model, by task condition and block. Error bars represent standard error of the mean.

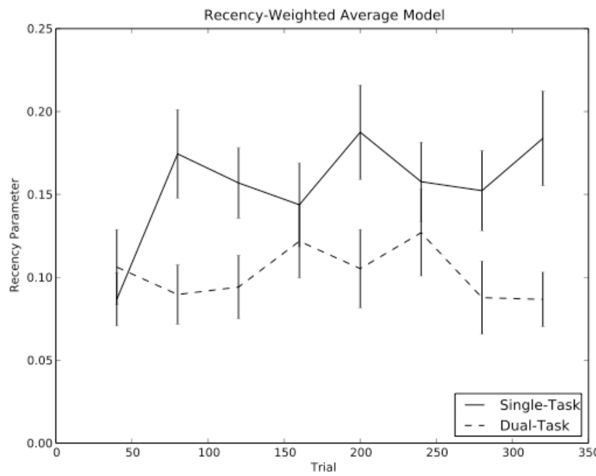

Models of the Two Prediction Strategies

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 WSLS usage, we fit a one-parameter WSLS model to participants' choices constraining $P(shift|lose)=P(stay|win)$, hypothesizing that ST participants would be better fit by this model than DT participants. To measure EM usage, we fit a model whose probability of making a green prediction at trial t is equal to the proportion of green outcomes in a growing historical window of outcomes leading up to time t. This zero-parameter model 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 participant's choice data using maximum likelihood estimation, and compared the relative goodness-of-fit of the

[2] Intuitively, participants' usage of WSLS should be reflected by a recency parameter values near one. As can be seen in Figure 2, ST participants' recency parameters were much lower than one. However, fitting the recency-weighted averaging model to simulated WSLS choice behavior reveals that a fairly deterministic WSLS strategy defined by $P(shift|lose)= P(stay|win)= 0.8$ results in recency parameter estimates near 0.3. Thus a participant that relies heavily on WSLS will be best fit by recency parameter values in this range.

models using the Akaike Information Criteria (AIC; Akaike, 1974) to correct for differing numbers of parameters. A lower AIC value for a model indicates a better fit. Intuitively, AIC_{EM} - AIC_{WSLS} yields a measure of how much better a participant is described by WSLS compared to EM.

Figure 3.

Comparison of model goodness-of-fit between WSLS and EM models. Average AIC value differences, calculated as $AIC_{EM} - AIC_{WSLS}$ using best-fitting parameter values for each block of each participants' choices. Error bars represent standard error of the mean.

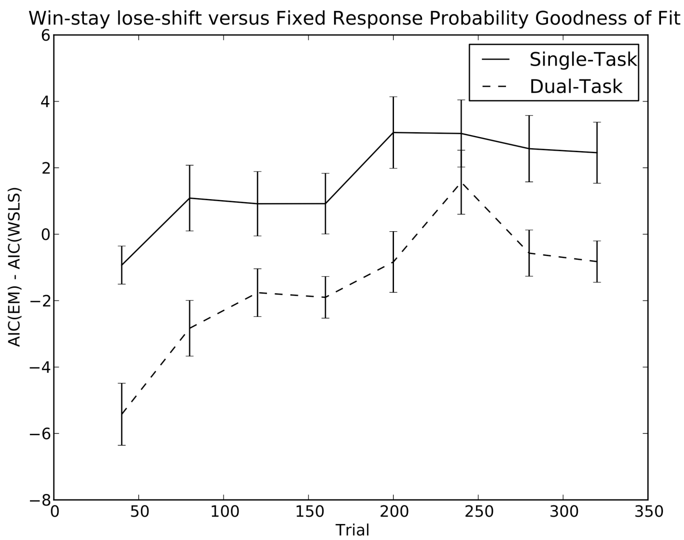

We predicted that the WSLS model would do a better job (relative to EM) of describing ST participants than DT participants. Figure 3 depicts the relative goodness-of-fit between the two models, for each condition across the 8 blocks. Indeed, the likelihood ratios reveal that ST participants were better described by the WSLS model than the responses of DT participants, and conversely, DT participants were better described by the EM model. A 2 (task condition) x8 (trial block) ANOVA revealed a significant main effect of task condition, $F(1,121)=11.97$, $p<.001, \eta_p^2=.099$, a main effect of trial block, $F(1,121)=36.63, p<.001, \eta_p^2=.303$, and no significant interaction between task condition and trial block, $F(1,121)=.62, p=.43, \eta_p^2=.005$. The main effect of task condition suggests that working memory load influenced the prediction strategies utilized by decision-makers.

Table 1

Correlations between self-reports of prediction strategy and model goodness-of-fit. Note that smaller AIC values indicates a better model fit, and larger self-report values indicate greater endorsement of a strategy.

	Correlation between Self-Report and Model Goodness-of-fit (AIC)	
Condition	WSLS	EM
Single Task	r(68) = -.36, p<.001	r(68) = .15, p = .23
Dual Task	r(51) = -.02, p = .90	*r*(51) = .06, p = .64

Strategy Self-Reports

As discussed above, we predicted that if WSLS is an explicit strategy and EM is not, that in a single-task setting people should be able to predict their use of WSLS, but not EM. In a dual-task setting, identifying any strategy explicitly may be impaired by the WM load. To examine this prediction, we related participants' offline endorsement of the strategies that were described in the questionnaire to the model fits for each strategy. 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 EM). Table 1 reports correlations between endorsement and overall model goodness-of-fit (because AIC is used, negative correlations indicate predictive relationships). As hypothesized, neither condition's endorsement of EM correlated with model goodness-of-fit. However, self-reports of ST participants accurately predict WSLS usage while those of DT participants do not[3]. Regressing AIC_{WSLS} against WM load and WSLS endorsement confirmed this interaction between WM load and strategy endorsement (β=8.80,*SE*=3.48,p<.05).

Discussion

We investigated the effect of a concurrent WM load—which is believed to disrupt explicit (or declarative) learning on individuals' prediction behavior in a sequentially independent series of outcomes. While both Single-Task (ST) and Dual-Task (DT) participants appeared to probability-match to the same extent globally, model-based analyses suggest that WM load drove local choice behavior away from win-stay-lose-shift (WSLS)—under which, choices rely on the most recent outcome—and towards expectation matching (EM), which relies on outcomes integrated over a longer historical window. The depletion of WM resources resulted in reliance upon a larger window of past information, which is somewhat counterintuitive but also consistent with contemporary distinctions of explicit/implicit processing (e.g., Ashby & Maddox, 2005; Henke, 2010).

Our results are interesting in the context of previous dual-task studies. Foerde and colleagues (2006) found that a concurrent WM load during probabilistic classification learning impaired participants' acquisition of explicit associations between perceptual cues and outcomes, although these participants evidenced implicit learning of cue-outcome contingencies. Further, they were unable to flexibly apply cue knowledge offline, suggesting that WM load engenders the use of implicit learning systems. Likewise, the present study found a discrepancy between ST and DT groups' ability to explicitly encode knowledge about outcome frequencies and identify strategies they employed. These discrepancies, taken in conjunction with model-based analyses identifying choice strategies, support our view that WSLS and EM are characteristic of explicit

---

[3] Maximizing is another type of explicit strategy, and hence, even though it was rare (though more likely in later blocks), we predicted a significant correlation between maximizing strategy use and self-reports. As a proxy for maximizing behavior, we fit a modified EM model to the second half of participants' choice data, which predicts always choosing the higher-frequency outcome in its history. Our prediction was confirmed: AIC_{MAX} significantly correlated with participants' rankings of the maximization strategy, r(115)=-.23, p=.01. These correlations did not differ significantly between the ST group, r(67)=-.22, p=.06, and the DT group, r(48)= -.23, p=.10 [z=-.02, p=.98.].

and implicit modes of operation respectively[4]. Whether these two strategies depend on two distinct neural systems or a single system operating in disparate modes is matter of future investigation.

It is worth noting that the EM model fit to participants' data does not distinguish between the EM strategy described earlier and a simpler strategy whereby the decision-maker makes a prediction based on a randomly retrieved choice among outcomes seen so far. If it was the case that this alternate strategy was utilized by DT participants,

Previous work has examined the effects of a verbal WM load on prediction behavior in a task similar to ours (Wolford et al., 2004), finding that memory load rendered decision-makers less likely to probability-match and more likely to maximize. While we found no significant differences in maximizing behavior between ST and DT participants, there are two important differences between these studies that could explain the divergent results: our secondary task was likely more difficult, and our outcome base rates were closer to equiprobable (65% compared to 75%) which is unlikely to foster maximization.

It has been proposed that in probabilistic and sequentially independent tasks decision-makers may form incorrect beliefs about the outcome-generating process, which, in turn, inform prediction behavior (Green, Benson, Kersten, & Schrater, 2010). Under one proposal, molar-level probability-matching may result from a search for deterministic patterns in an attempt to achieve higher accuracy (Gaissmaier & Schooler, 2008; Wolford, Miller & Gazzaniga, 2000). As this work suggests pattern search is WM-intensive, it is unlikely that DT participants would have engaged in search. The present results do not rule out the possibility that ST participants could have engaged in a pattern search. However, this strategy is not readily identifiable with the descriptive modeling approach taken here. Future work is needed to develop models that can identify a pattern-search strategy, and further, to investigate the effect of a WM memory load on this strategy.

---

[4] Though we assume an implicit vs. explicit distinction to account for our results, another possibility is that the strategies differ in their complexity, and thus, the amount of working memory required to reflect upon and identify them later during the strategy reports. Our results cannot distinguish between these accounts, but future studies may vary these factors independently to identify the degree to which complexity accounts for the results presented here. We thank an anonymous reviewer for pointing out this possibility.

References

Akaike, H. (1974). A new look at the statistical model identification. *IEEE Transactions on Automatic Control*, *19*(6), 716-723.

Ashby, F. G., Alfonso-Reese, L. A., Turken, A. U., & Waldron, E. M. (1998). A neuropsychological theory of multiple systems in category learning. *Psychological Review*, *105*(3), 442-481.

Ashby, F. G., & Maddox, W. T. (2005). Human Category Learning. *Annual Review of Psychology*, *56*(1), 149-178.

Fantino, E., & Esfandiari, A. (2002). Probability matching: encouraging optimal responding in humans. *Canadian Journal of Experimental Psychology, 56*(1), 58-63.

Filoteo, J. V., Lauritzen, J. S., & Maddox, W. T. (2010). Removing the Frontal Lobes: The Effects of Engaging Executive Functions on Perceptual Category Learning. *Psychological Science, 21*(3), 415-423.

Foerde, K., Knowlton, B. J., & Poldrack, R. A. (2006). Modulation of competing memory systems by distraction. *Proceedings of the National Academy of Sciences*, *103*(31), 11778-11783.

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

Green, C. S., Benson, C., Kersten, D., & Schrater, P. (2010). Alterations in choice behavior by manipulations of world model. *Proceedings of the National Academy of Sciences*, *107*(37), 16401 -16406.

Henke, K. (2010). A model for memory systems based on processing modes rather than consciousness. *Nature Reviews Neuroscience*, *11*(7), 523-532.

Herrnstein, R. J., Rachlin, H., & Laibson, D. I. (2000). *The Matching Law*. Cambridge, MA: Harvard University Press.

Hinson, J. M., & Staddon, J. E. R. (1983). Hill-climbing by pigeons. *Journal of the Experimental Analysis of Behavior*, *39*(1), 25-47.

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.

Rescorla, R., & Wagner, A. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In *Classical Conditioning II: Current Research and Theory* (pp. 99, 64). New York: Appleton-Century-Crofts.

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.

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.

Vulkan, N. (2000). An Economist's Perspective on Probability-matching. *Journal of Economic Surveys*, *14*(1), 101-118.

Wolford, G., Miller, M. B., & Gazzaniga, M. (2000). The Left Hemisphere's Role in Hypothesis Formation. *Journal of Neuroscience, 20*(6), RC64.

Acknowledgements

The authors wish to thank Brad Love, Alex Huk, and Russ Poldrack for their helpful discussions. The authors also thank J. Grant Loomis for his assistance with data collection and two anonymous reviewers for their helpful comments.