Betting on the Unexpected: The Effect of Expectation Matching on Choice Strategies in a Binary Choice Task

by

Greta James

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Master of Arts
in
Psychology

Waterloo, Ontario, Canada, 2012

Author's Declaration

I hereby declare that I am the sole author of this thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

____

Greta James

Abstract

Probability matching is the tendency to predict outcomes in accordance with their actual contingencies in a binary choice task. It is,however, a suboptimal response if the aim is to maximize correct predictions. I review two theories that attempt to explain why probability matching occurs: the pattern-search hypothesis and dual-systems theory. These theories are tested in two studies which suggest that dual-systems theory provides a better account of probability matching behavior. Studies 3, 4, and 5 then provide evidence for an extension of the dual-systems theory, called expectation matching, which is intended to explain why probability matching is the intuitive response to a binary choice problem.

Acknowledgements

This research would not have been possible without the help and extreme patience of my faculty supervisor, Derek Koehler. I would also like to thank my readers, Colin MacLeod and Daniel Smilek, for the time they took to read and comment on this thesis. It has been greatly improved by their efforts. Finally, I would like to thank Bill Eickmeier for the time he spent helping me program many of the experiments reported in this thesis.

The experiments and results reported here have also been published in the following articles:

James, G., & Koehler, D. J. (2011). Banking on a bad bet: Probability matching in risky choice is linked to expectation generation.*Psychological Science, 22*(6), 707-707-711.doi:10.1177/0956797611407933

Koehler, D. J., & James, G. (2009). Probability matching in choice under uncertainty: Intuition versus deliberation. *Cognition, 113*(1), 123-127. doi:10.1016/j.cognition.2009.07.003

Koehler, D. J., & James, G. (2010). Probability matching and strategy availability. *Memory & Cognition, 38*(6), 667-676. doi:10.3758/MC.38.6.667

Table of Contents

List of Figures

Introduction

Most prominent models of decision making share the assumption that people will choose a gamble offering a higher probability of winning over and otherwise equivalent (and equivalently framed) gamble that offers a lower probability of winning. Indeed, this assumption underlies the very notion that people are rational decision makers. However, since the 1950s researchers have been investigating a the tendency of some decision makers to choose a lower probability gamble over a higher probability gamble when engaged in a phenomenon known as probability matching. Probability matching cannot readily be accommodated by even the most influential models of decision making, such as prospect theory (Kahneman&Tversky, 1979). Furthermore, probability matching is tends not to occur in other, presumably less rational, mammals (Bullock &Bitterman, 1962; Lee et al.,2004; Wilson, 1960). Even young children under the age of 5 (Derks&Paclisanu, 1967)tend to behave more rationally when engaged in tasks which elicit probability matching in older children and adults. As such, why adult humans engage in the puzzle of probability matching has been an intriguing puzzle for researchers for the last 60 years because it presents such a fundamental challenge to most decision-making theories.

But what is probability matching? Probability matching occurs when participants are given a simple choice between two gambles, one with a greater than 50% chance of winning a reward (usually money) and the other with a less than 50% chance of winning a reward. It is hard to imagine that anyone would have trouble selecting the optimal option when presented with such a choice. However, if you present this choice as a series of gambles, many adult humans will indeed choose the less than 50% odds some of the time.

To illustrate, imagine a 10-sided die on which 7 sides are painted green and 3 sides are painted red. Suppose I tell you that I am going to roll this die 10 times in a row and before each roll I will ask you to predict the outcome. For every roll that you predict correctly, you will win $1. When playing such a game, people tend to choose one of two strategies. The rational strategy, called maximizing, is to choose the more probably option on every trial (i.e.: guessing all green). This is the strategy predicted by most theories of decision making. Unfortunately, a second strategy is also popular. This strategy, known as probability matching, is to choose each color in accordance with its probability; in other words, choose 7 greens and 3 reds. Of course, this strategy is sub-optimal, because choosing any reds at all is equivalent to choosing the 30% odds of winning a dollar over the 70% odds of winning a dollar. Despite the obvious benefits of maximizing, over the years many studies have demonstrated probability matching(seeGrant, Hake,&Hornseth, 1951, for seminal work on probability matching and Vulkan, 2000, for a review).

Probability matching is not only common, it is also quite resistant to changes in experimental design. Participants will probability match in any situation involving two choices in which the outcomes are randomly determined, but one choice happens with greater frequency than the other. For example, binary choice tasks can take a variety of forms, including word problems(West &Stanovich, 2003), die outcomes(Gal & Barron, 1996; West &Stanovich, 2003), marbles drawn from a bag (Gal & Barron, 1996), stimuli such as lights presented on a computer screen (Yellott, 1969), and decks of cards (Brackbillet al., 1962). Participants may learn the outcome probabilities through description (as described in the die problem above) (West &Stanovich, 2003), or they may learn them through experiencing a number of outcomes (Tversky& Edwards, 1966). Participants may receive feedback as to the accuracy of their guesses (Newell &Rakow, 2007), or they may make all guesses without any feedback (West &Stanovich, 2003). Participants may be paid for correct guesses (Shanks et al., 2002), or only imagine they are being paid for correct guesses (Gal & Barron, 1996). In all of these cases, probability matching behavior persists to varying degrees.

Given the prevalence, resilience, and importance of probability matching behavior, it is remarkable that after more than 50 years of research we still are not able to explain why it occurs. Current theories fail to accommodate all of the available evidence on the phenomenon. In particular, I will compare and contrast two alternative accounts of probability matching behavior: the pattern-search hypothesis and the dual-systems hypothesis. I will begin by reviewing current research in support of these theories and then present two experiments designed to compare and contrast them. Finally, I will present an addition to the dual-systems account of probability matching which improves its explanatory power. The final experiments presented in this thesis are designed to test this addition.

PatternSearch or DualSystems?

The patternsearch and dualsystems hypotheses provide two competing explanations of probability matching behavior. The literature provides evidence congruent with one theory or the other but neither theory has been tested conclusively.

The patternsearch hypothesis, argues that probability matching behavior is an adaptive response to environments in which outcomes potentially follow predictable patterns. Such an environment "primes" us to look for patterns in sequences that appear non-random in order to eventually achieve perfect predictive accuracy. Of course, this strategy only makes sense if participants believe that a sequence actually contains a pattern, in spite of being told that the sequence is random. There is a convincing collection of research (see Falk &Konold, 1997, for review) demonstrating that people are poor at recognizing and generating random sequences. In fact, one must create a decidedly non-random sequence involving large amounts of alternation for the sequence to appear random to the typical human observer (Lopes & Oden, 1987). The patternsearch hypothesis argues that, because sequences appear non-random to participants, they engage in a pattern match activity which involves generating potential pattern sequences that share the outcome probabilities of the chance mechanism. The result is what appears to be probability matching.

One problem with the patternsearch hypothesis to date is that very few studies have actually tested it. Most studies offer evidence that is congruent with the patternsearch hypothesis, but could also easily accommodate many other accounts. For example, Wolford et al. (2004) found that taxing working memory led to more maximizing behavior. They argued that this resulted from an inability to search for patterns under working memory load. In reality, however, all this finding really demonstrates is that probability matching requires more working memory than does maximizing. This can easily be accommodated by a need to monitor past responses in order to produce a sequence with the required contingencies: It does not necessitate a search for patterns.

In another example of congruent evidence, researchers demonstrated that if a sequence was rigged such that participants perfectly predicted the "random" sequence (i.e.the sequence was perfectly response dependent), they would later report using elaborate patterns to achieve perfect predictive accuracy (Unturbe&Corominas, 2007;Yellott, 1969). This was taken as evidence that participants are searching for patterns in all binary choice tasks. At face value, this provides compelling evidence, but on reflection it is hard to imagine participants reporting anything else. If one finds that they are demonstrating nearly impossible accuracy at predicting a random event, they are really only left with two conclusions: Either the event is not actually random (at which point it is logical to begin searching for patterns) or the participant is experiencing ESP. Given these alternatives, it is not surprising that the majority of subjects opt for the patterns. This does not suggest,however, that they were initially searching for patterns.

Neurological evidence has also been used to support the pattern matching hypothesis. Research has demonstrated that the left side of the brain tends to favor a matching strategy while the right side of the brain prefers to maximize (Miller et al., 2005; Wolford et al., 2000). Wolford et al.(2000)demonstrated that in patients with a left hemisphere lesion, maximizing prevails. They argued that the left hemisphere is important for searching for patterns and without a properly functioning left hemisphere participants revert to a maximizing strategy. However, they failed to demonstrate directly that the left hemisphere's function in probability matching is related to a search for patterns. The left hemisphere's role could encompassany number of functionsincluding, for example, the ability to sequence two alternative responses.

Finally, Gaissmaier and Schooler (2008)reported that those participants who used a matching strategy in a standard probability learning task (with serially independent outcomes) were more likely to identify and exploit a pattern when they encountered a sequence that was non-random. This directly links probability matching with improved pattern recognition, but still allows for alternative explanations. For example, carrying out a probability matching strategy may require that more attention be allocated to the task and may, therefore, make these participants more likely to notice patterns by virtue of increased attention.

The above studies are typical of the support provided for the pattern matching hypothesis in that they fail to directly test the patternsearch hypothesis and instead provide evidence that is congruent with the hypothesis, but may also be congruent with many other explanations of probability matching behavior. In general, very few studies have actually systematically varied pattern information to determine what effect this might have on probability matching behavior. One exception is a study by Wolford et al.2004in which pattern information in sequences was increased by including more alternationof the two possible outcomes. In the die problem presented earlier, this would mean participants still saw 70% green outcomes and 30% red outcomes, but the number of times a green outcome was followed by a red outcome and vice versa was increased. Ironically, such a sequence tends to appear more random to participants (Lopes & Oden, 1987) and Wolford et al. found that it led to increased maximizing. In the case of increased alternation, we can conclude that if you make a sequence more patterned, participants will cease looking for patterns and start maximizing! Nevertheless, this study provides strong support for the notion that random sequences appear non-random and this is in some way related to a probability matching strategy.

An alternative explanation of probability matching behavior comes from dualsystems theory (Kahneman& Frederick, 2005). This theory suggests that to correctly solve a problem one must be able to access all available solutions and then choose the best solution among them. In general, many problems tend to lend themselves to the generation of some solutions and not others. These easily generated solutions are referred to as intuitive solutions (system 1), while more difficult to generate solutions are labeled deliberative solutions (system 2). The latter can only be arrived at if participants are able to override the more intuitive solution and both generate and realize the superiority of the more deliberative solution. In the case of binary choice tasks, dualsystems theory would argue that probability matching is the intuitive strategyand maximizing is the deliberative strategy. Probability matching is selected as the intuitive strategy because it is (a) erroneous and (b) highly available. Maximizing could require deliberation because it is not easily generated as a strategy or because it is not often recognized as superior to probability matching, or both.

To date, there is only a small amount of research which may provide support for a dualsystems account of probability matching. Like the research examining the patternsearch hypothesis, much of this research does not provide direct support for the dualsystems account, but rather provides evidence that is congruent with the theory. For example, manipulations that can be interpreted as encouraging deliberation, such as instructing participants to recommend a strategy to another person (Fatino&Esfandiari, 2002) or to think like a statistician (Kogler&Kuhberger, 2007), have been found to increase maximizing behavior. Furthermore, maximizing is also more common in situations that encourage consideration of alternative strategies, such as comparing participants' performance to that which could have been achieved using an optimal strategy (Shanks, Tunnery,& McCarthy, 2002). Stanovichand West (2008) found that those higher in cognitive ability, and thereby presumably more efficient in deliberative reasoning, are also more likely to maximize. Finally, the notion that probability matching is an intuitive mistake also accommodates findings such as increased payoff leading to greater maximization (Brackbillet al., 1962; Shanks et al.,2002). Under increased stakes, people may be more likely to deliberate and therefore also more likely to maximize. Such a finding is more difficult to accommodate under thepatternsearch hypothesis as increased incentive would presumably lead to a more determined patternsearch and more probability matching.

Although Dual systems theory may provide a convincing account of this collection of evidence, these findings are really only loosely linked to deliberative reasoning. Concrete tests of dual systems theory's application to probability matching are even fewer than those reported for the pattern search hypothesis.

To this point, then, adjudicating between the patternsearch hypothesis and the dualsystems account of probability matching awaits further research. Neither account is conclusively supported by what research has already been done. Fortunately, both accounts make considerably different and easily testable predictions about a common difference in the existing probability matching literature. This difference, which I will refer to as the method discrepancy refers two different ways presenting the dual task used to test for probability matching. In *choice-based studies*, participants are classified as maximizers or probability matchers based on a summary of the individual choices that they made on a binary choice task. By contrast, *strategy-based studies* explicitly provide participants with the maximizing and probability matching strategies and the participant simply indicates which strategy he or she thinks is optimal.

To illustrate each method more clearly, consider the die problem mentioned above. In a choice-based study, participants would be shown the die and asked to make 10 different predictions in which they guessed either red or green. These predictions would later be added up by the researcher and those participants that guessed roughly 7 green and 3 reds would be classified as probability matchers whereas those who guessed all green would be classified as maximizers. In a strategy-based study, the die problem would be described to participants as well as both strategies and they would be asked to indicate which strategy they thought was superior. Choice-based studies (e.g., Edwards, 1961; Tversky& Edwards, 1966; West &Stanovich, 2003, Exp. 3) tend to report lower levels of maximizing than do strategy-based studies (e.g., Gal & Baron, 1996; Stanovich& West, 2008; West &Stanovich, 2003, Exp. 1–2). This method discrepancy is of particular interest because it lends itself to two distinctexplanations, one provided by the patternsearch hypothesis and the other by a dualsystems account of probability matching.

The patternsearch hypothesis would suggest that one reason that probability matching might be more common in choice-based problems is because there is pattern information in such problems, but not in the strategy-based problems. If participants probability match because they are searching for or testing a pattern, they must see independent outcomes of the event and be given the opportunity to produce independent predictions. Without such experience, there is no opportunity to search for patterns or achieve perfect predictive accuracy by employing a pattern.

Dualsystems theory provides a different prediction with respect to the method discrepancy. According to this account, the availability of pattern information should make no difference to rates of probability matching because it does not provide information about competing strategies or make competing strategies easier to evaluate. Strategy-based studies lead to more maximizing because they make all strategies equally available, thus removing one obstacle to overcoming the intuitive probability matching response. In other words, in strategy-based studies, participants do not need to generate all possible strategies; they only need to be able to evaluate which one is better. By contrast, choice-based strategies provide no such advantage.

Thus, the patternsearch hypothesis holds that strategy-based studies lead to more maximizing because participants do not observe a sequence of outcomes and, thus, no pattern information is available. If the dualsystems hypothesis is correct, varying the accessibility of pattern information should make no difference to the degree of probability matching behavior observed. Directly comparing situations in which participants have the opportunity to observe and respond using pattern information with situations where this is not an option should provide an excellent test of these two hypotheses. Study 1 examines the effect of pattern information on probability matching behavior. If participants are indeed searching for patterns, then removing all pattern information from the problem should make participants more likely to maximize.

Study 1

Study 1 investigates the role that access to pattern information plays in participants' strategy choice. Participants may or may not receive pattern information would learning about the underlying contingencies and may or may not be able to include pattern information when making their responses. Those interested in looking for and testing patterns will be unable to do so if pattern information is unavailable and so this manipulation should decrease rates of probability matching if the pattern search hypothesis is correct.

METHOD

Participants

Participants were 120 undergraduate students (53 female) recruited from a campus student centre. They were told that they would receive up to $10 for their participation depending on their performance.

Procedure

The computer-based choice task was described as a game in which participants were to guess the colors of marbles that were to be drawn from a bag containing a mix of red and green marbles. The task consisted of a learning phase followed by a test phase. In the serial learning condition, participants saw 40 marbles drawn, one at a time, from the bag; in total they saw 30 green and 10 red marbles[1] drawn in a randomized order. This condition allowed participants to search for patterns should they be inclined to do so, although in fact the outcomes were serially independent. In the aggregate learning condition, participants were told that a total of 30 green marbles and 10 red marbles had been drawn from the bag, but they were not presented with trial

[1] For ease of exposition, green will be referred to as the dominant color in the task; in fact, the dominant color was counterbalanced across participants.

by trial outcomes and so had no opportunity to observe potential patterns. In both learning conditions, participants were told that each of the 40 marbles had been drawn randomly, with replacement, from the bag.

In the test phase, participants were told that 20 more marbles would be drawn, with replacement, from the same bag, and that they would earn $0.50 each time they correctly guessed the color of a marble drawn from the bag. In the serial test condition, participants were asked to guess the color of each marble drawn, one at a time, but without feedback regarding the color that was actually drawn on each trial. The serial test condition allowed participants to order their guesses to follow a pattern, should they choose to do so. In the aggregate test condition, participants were asked to indicate how many times, across the 20 draws, they would guess red, and how many times they would guess green. In this condition, even if participants suspected that that the outcomes might follow some sort of pattern, they had no way to exploit that pattern in making their responses.

There were 30 participants assigned to each of the four cells resulting from this 2 (learning: serial versus aggregate) by 2 (test: serial versus aggregate) design. After finishing the choice task, participants completed a questionnaire in which, among other items, they estimated the overall proportion of red and green marbles in the bag and then evaluated alternative choice strategies. Specifically, a number of strategies were described, including matching and maximizing, which might have been used during a game similar to the one that the participant had played. In the case of this question, it was known that the bag in question contained 70% green marbles. Each strategy was described as having been used by a different hypothetical player and participants were asked to rank the players in terms of their expected payoffs. After they completed the questionnaire, participants were paid on the basis of their guesses and associated outcomes on the task.

RESULTS

Data from 8 participants who made more red than green guesses, or who mistakenly indicated red to be the more probable color in the follow-up questionnaire, were excluded from further analysis, as were data from 10 additional participants who failed to complete the entire study.

Overall, participants were largely accurate in their estimates of the probability of drawing a green marble (mean = 70%; SD = 8%).

An index of a participant's tendency to probability match versus maximize was derived by comparing the proportion of times that the participant chose to guess green (C) to his or her estimate of the proportion of green marbles in the bag (E, taken from the follow-up questionnaire) as follows:

$$\text{choice index} = (C - E)/(1 - E)$$

This index adjusts for differing perceptions of the probability of drawing a green marble from the bag. Probability matching gives rise to a score of 0 (because C = E) whereasmaximizing results in a score of 1 (because C = 1). Negative scores indicate choice probabilities that are lower than estimated probabilities (e.g., participants who made an equal number of green and red guesses despite estimating green to be more probable).

Across all conditions, many more participants engaged in probability matching (n = 46 participants with an index value of 0) than in maximizing (n = 14).Choices of the remaining participants did not fall in either category by these strict definitions, but the mean value of the index (0.12) and its distribution (see Fig. 1) underscore the appeal of strategies falling close to that expected by probability matching. An analysis of variance on the choice index measure indicated no significant main effect nor interaction between the two experimental manipulations, learning $F(1,98) = 0.003$,MSE=0.111, $p = 0.96$; test $F(1,98) = 0.45$, MSE=0.004, $p = 0.51$;learning by test $F(1,98) = 1.80$, MSE=0.131, $p = 0.18$. There was no evidence that matching was more prevalent when potential pattern information could be identified or exploited.

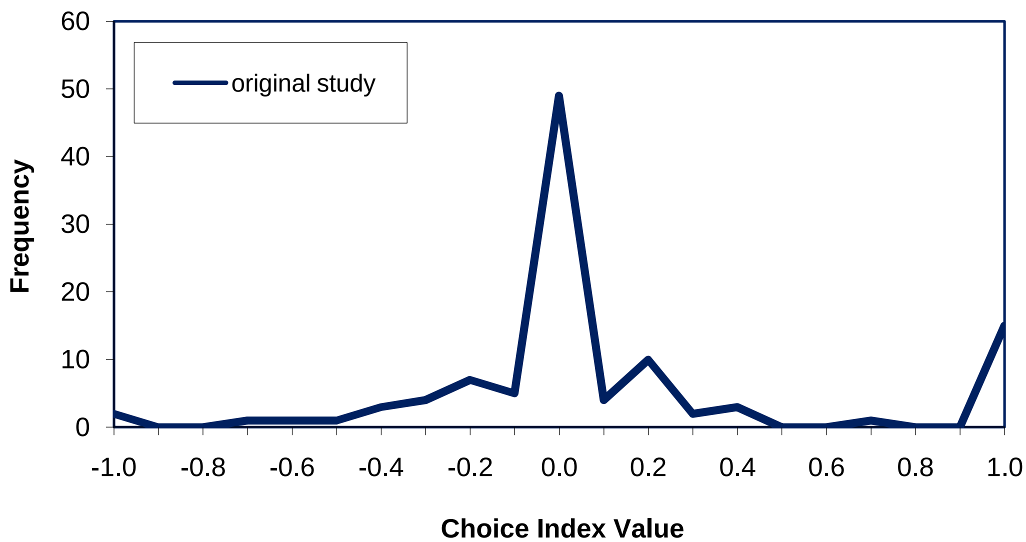

Fig 1*.Distribution of choice index value across participants in study 1. An index value of 0 indicates probability matching; a value of 1 indicates maximizing.*

Responses to the strategy question were categorized in terms of whether the maximizing strategy was ranked higher or lower than the matching strategy with respect to expected payoffs. When explicitly presented with both strategies and encouraged to compare them, many participants (n = 50) ranked maximizing as the better strategy; a slightly larger number (n = 52) ranked probability matching as better. The proportion of participants who ranked maximizing as superior on the questionnaire was significantly higher than the small proportion who actually used the maximizing strategy in the choice task, χ^2 (1, n =102) = 12.48, *p*<.001. Over 40% of participants (n = 37) who used a non-maximizing strategy during the choice task switched to endorse a maximizing strategy on the questionnaire, including 18 of 46 (39%) classified as strict probability matchers on the choice task. Of those participants who maximized on the choice task, only one (7%) switched to the matching strategy.

DISCUSSION

The patternsearch hypothesis was not supported: Participants showed no tendency to switch from probability matching to maximizing under conditions in which patterns could not be identified or exploited. Instead, probability matching was consistently more common than maximizing. While it is still possible that probability matching results from more of an implicit search for patterns, work such as that of Wolford et al. (2004), in which participants viewed a sequence which appeared more random and showed less probability matching as a result, suggests that probability matching is sensitive to the characteristics of the pattern and consequently that it should be possible to alter.

These findings suggest that the difference between choice-based and strategy-based studies is more likely explained by the increased availability of both strategies, as predicted by dualsystems theory, as opposed to the existence or absence of pattern information. Notably, a substantial proportion of the participants in this study who engaged in probability matching on the choice task later acknowledged the superiority of maximizing when both strategies were explicitly described for comparison. This suggests that both strategies were not available to participants when they were choosing how to respond to the problem. One possible reason why people engage in probability matching, then, is that it springs readily to mind as a strategy whereas maximizing does not. This is congruent with the ideas expressed by dualsystems theory. Recall that, according to this theory, probability matching would be the intuitive strategy, and it must be overridden to arrive at the maximizing strategy.

It is also possible that those participants that rated maximizing as superior to matching, but still matched on the choice task, may have known all along that maximizing was superior, but chose to match for some other reason. Boredom (Goodnow, 1955), risk aversion (Goodnow, 1955), and a possible increase in utility of correctly predicting the infrequent response (Brackbill& Bravos, 1962) are all possible reasons. It is unlikely that the utility of predicting the least frequent response is responsible, as in half of the conditions subjects were not able to directly predict the infrequent response, and these conditions did not differ significantly from those where they could make such a prediction. Boredom is also unlikely as the task was short in all cases and especially so for those that learned and tested in the aggregate condition. In any case, the best way to resolve this issue is to provide participants with the strategies *before* the choice problem and determine whether these participants maximize more readily than those who evaluate the strategy after the choice problem.

While Study 1 suggests that pattern information is not the reason for the method discrepancy, it remains an open question whether strategy availability *is* the reason. Study 2 willexamine this further by determining whether participants use the maximizing strategy more in a choice-based problem when the strategy is made available to them. Study 2 tests the hypothesis that making both the maximizing and the matching strategy more available should lead to more maximizing behavior.

Study 2

METHOD

Participants

Participants were students recruited from a campus student life centreor through psychology classes. Two variants of the study were conducted with similar methods and results, so they are reported together here. In the first variant, conducted in the lab, psychology undergraduate students (n = 60; 40 females) participated in exchange for course credit and played a low-stakes version of the game in which each correct guess paid $0.25. In the second variant, conducted in the campus student life center, university students (n = 98; 40 females) played a high-stakes version of the game in which they could win up to $10 as compensation for their participation; in this version of the game, each correct guess paid $1.00.

Apparatus

The task involved ten pairs of cups, with one red and one green cup in each pair. A $1Canadian coin[2] was hidden under either the red or green cup in each pair. For each participant, the hiding place of the coins was determined for each pair of cups by the roll of a ten-sided die with 7 green faces and 3 red faces (or vice versa; dominant color was counterbalanced across participants). Participants placed a black ring around either the red or green cup in each pair to indicate that they thought the coinwas hidden under that cup. The pairs of cups were placed on a table in a tent that hid the game and outcomes from passersby.

[2] For those that received course credit, a quarter was hidden under the cup rather than a $1 coin. For all further references to $1 coins, assume that the same procedure was also used with quarters for those in the variant with course credit as partial compensation.

Procedure

The experimenter explained the game to participants as they sat in front of the table with the pairs of cups. They were told that the coins had been hidden using the die, before they had entered the tent. Participants were told that they would receive $1for each correct guess. Afterward, half of the participants made all ten choices by placing the ten rings while half of the participants answered the following question:

> "In the game that you are going to play, loonies are randomly placed under red and green cups and you will be paid if you correctly guess the color of the cup the loonie is under. Before you start the game please consider the following two strategies:
>
> 1) You could choose green for all 10 sets of cups
>
> OR
>
> 2) You could choose green for seven sets of cups and red for three sets of cups.
>
> Which strategy, 1 or 2, do you think will win you the most money?"

This question served to make both strategies available to participants prior to playing the game and, as such, acted as a hint. Participants in the hint condition then made all ten choices in an identical way to those who had not received the hint. Upon completion of the choice task, but before the cups were lifted, all participants completed a brief questionnaire. This questionnaire included the hint question described above and also asked participants to indicate which strategy (a) their own strategy more closely resembled; (b) would be expected to earn more money; and (c) they would use if they were to play the game again. The cups were then turned over and participants learned how much money they had earned. Finally, participants completed the Cognitive Reflection Test (CRT; Frederick, 2005), a measure of chronic proneness to reliance on intuition or deliberation, and several other items including a self-rating of proficiency in mathematics and a question about the number of mathematics courses the participant had taken since high school.

RESULTS

The two variants of the study produced nearly identical results. Rates of maximizing in each condition were very similar across the two versions of the study. In a 2 (study variant: low- vs. high-stakes) by 2 (condition: hint vs. control) ANOVA with number of dominant color guesses as the dependent variable, study variant had no significant main effect nor was its interaction with condition statistically significant. Hence, we collapse over the study variable in subsequent analyses.

In the same ANOVA, condition (hint versus control) had a significant main effect on number of dominant color guesses, $F(1,154) = 7.17$, MSE =20.861, $p = .008$. In the control condition, participants chose the dominant color an average of 7.9 times out of 10; by contrast, in the hint condition, participants chose the dominant color 8.7 times out of 10. Figure 2a shows the full frequency distribution of this variable by condition. The proportion of participants engaging in strict maximizing in the no-hint condition was 35% (27 out of 78 participants). Among participants in the hint condition, the rate of strict maximization was significantly higher at 53% (42 out of 80 participants), χ^2 (1, n =158) = 5.14, $p = .023$. The hint manipulation also had a marginally significant impact on self-reported strategy use as elicited following the choice task, χ^2 (1,n =158) = 3.05, *p*=.081, with 47 of 80 participants (58%) in the hint condition stating that their strategy more closely resembled maximizing than matching, versus only 35 of 78 participants (45%) in the no-hint condition.

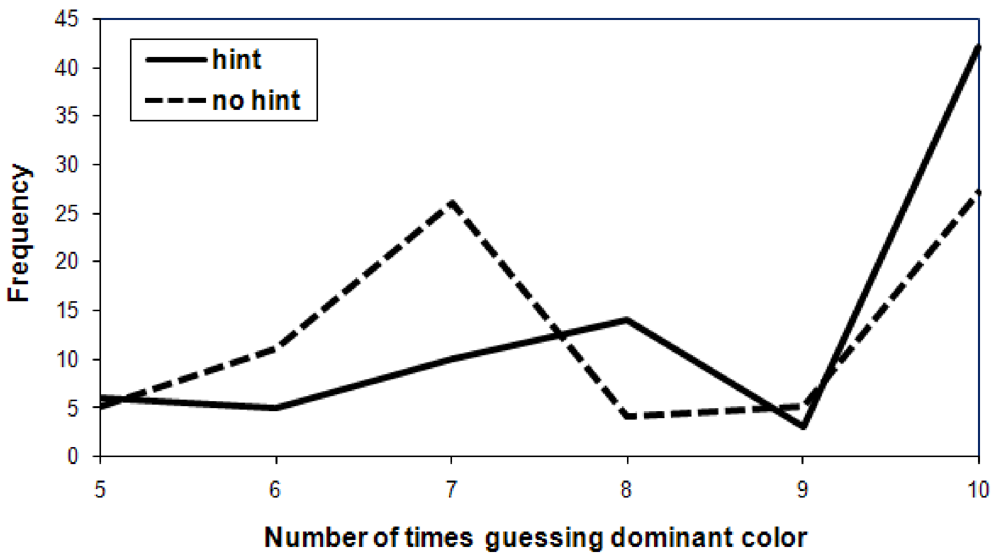

Fig. 2a.*Distribution of number of dominant color choices in Study 2, where the expected number of such choices is 7 given a matching strategy and 10 given a maximizing strategy, for the hint (solidline) and no-hint (dashed line) conditions.*

The hint manipulation did not have a significant effect on the other strategy questions asking which has the higher expected payoff and which the participants would use if they were to play the game again. On both measures, the maximizing strategy was selected by the majority of participants: 61 out of 80 participants in the hint condition and 62 out of 78 in the no-hint condition selected maximizing as having the higher expected payoff, and 51 of 79 participants in the hint condition (one person did not complete this item) and 53 of 78 participants in the no-hint condition said they would maximize if they were to play the game again. Thus both groups endorsed the maximizing strategy to the same extent when it was brought to their attention as an alternative to matching but, when this question was posed before the choice task, participants were more likely to use it in making their choices. Figure 2b shows, for each condition, the proportion of participants classified as matchers or maximizers on the choice task[3] as well as endorsement rates of the strategies on the three strategy comparison questions.

To clarify the impact of the hint manipulation on choice task performance, we can restrict our analysis to the subset of participants who endorsed maximization as the better strategy when presented with the strategy comparison question after completing the choice task.[4] This analysis provides an estimate of the influence of strategy availability, then, among the subset of participants who are able to identify the better strategy in a direct comparison. As noted above, the number of participants who chose maximizing as superior in the post-task strategy comparison was virtually identical in the hint and no-hint conditions (n = 61 and 62, respectively), indicating that participants in both groups were equally able to recognize the better strategy in the direct comparison. But when this strategy was not readily available at the time that the choice task was performed, in the no-hint condition, the rate of strict maximization was much lower (27 out of 62 participants, or 44%, maximized) than when it was readily available, in the hint condition (42 out of 61 participants, or 69%, maximized), χ^2 (1, n =123) = 7.99, p = .005.

Next we turn to scores on the CRT, which consists of three mathematical problems, each of which gives rise to an initial intuitive answer that can clearly be identified as incorrect when subjected to further scrutiny (Frederick, 2005). Thus, the number of correct answers on the CRT

---

[3] Participants were classified as maximizers if they guessed green 9 or more times and matchers if they guessed green 6, 7 or 8 times. This coincides closely with their self-ratings of their own strategy use.

[4] Results of this analysis do not change when we consider hint-condition participants' pre-choice-task responses to this question instead of their post-task responses.

can be taken as an index of the individual's ability to override their initial intuition in favor of a more deliberative approach to the problem that yields the correct solution. CRT scores correlated positively with number of dominant color guesses, $r = .36$, $p < .001$. CRT remained

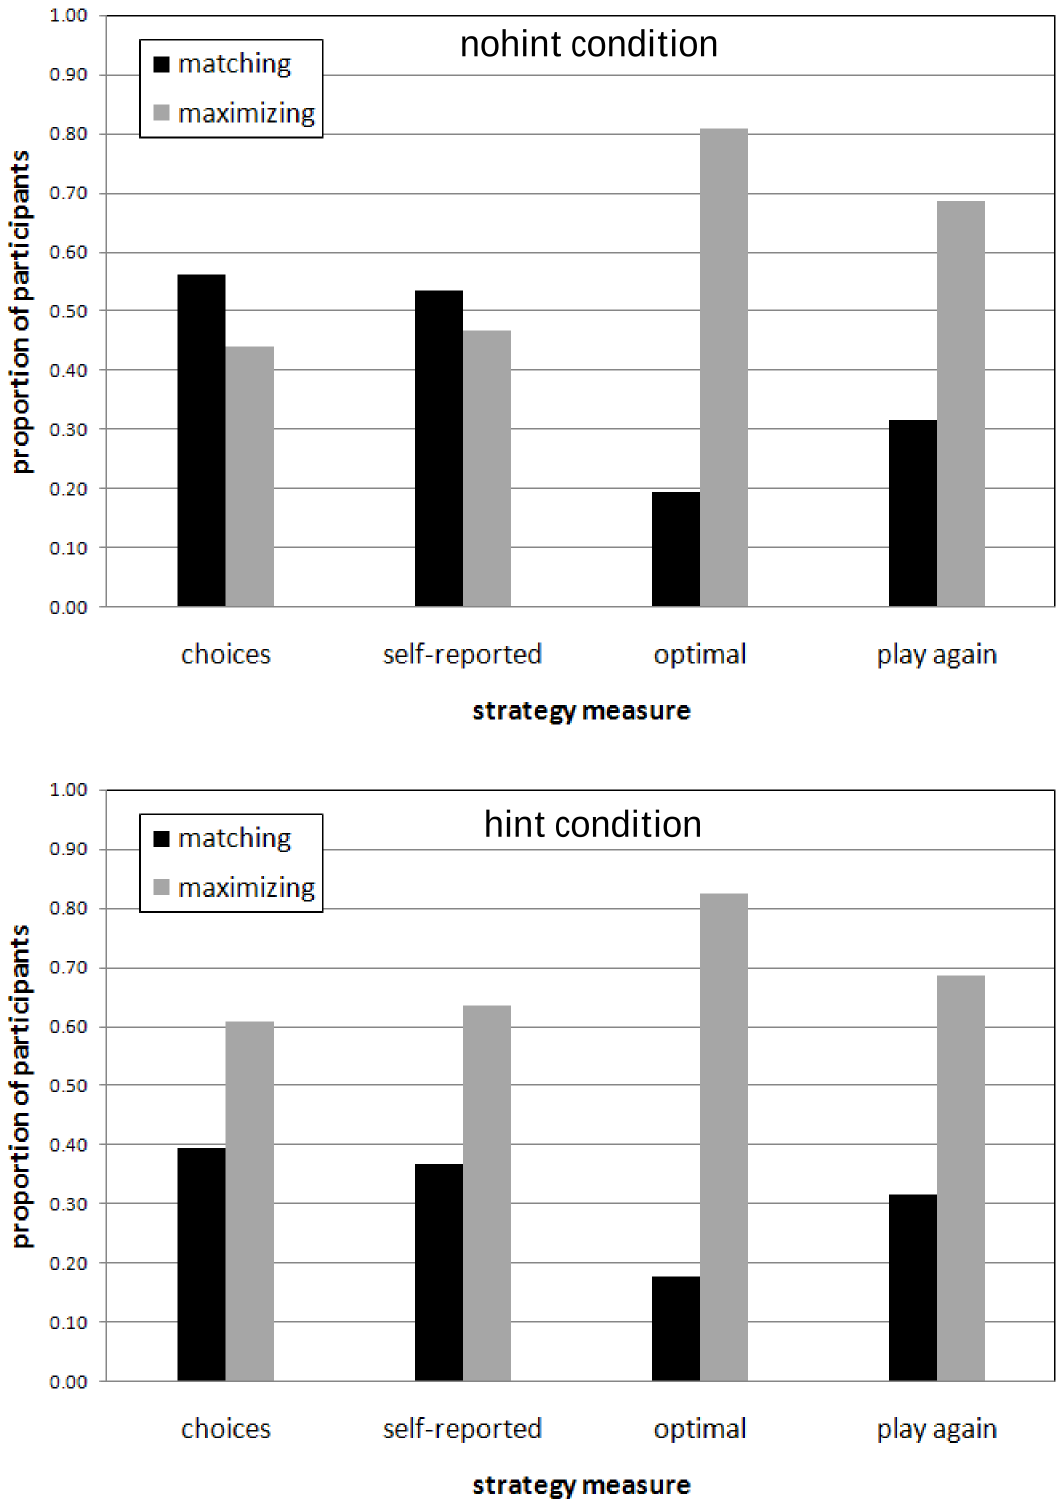

Fig.2b: *Strategy use and endorsement rates in Study 2, by condition.*

statistically significant when entered simultaneously as a predictor with number of math courses taken, providing some indication that it accounts for something other than simple ability at math.

Finally, the correlations between CRT and the number of dominant color guesses on the choice task were computed separately for the hint and the no-hint conditions. Althoughthe correlations remained significant when computed within each condition, they tended to be stronger in the hint than in the no-hint condition: In the hint condition, the correlation with number of dominant color choices was .43 for CRT; in the no-hint condition, the corresponding correlation was .28 for CRT.Compared to participants prone toward reliance on intuition, those prone toward reliance on deliberation apparently benefited more from the hint. Figure 6 displays this pattern of results for the CRT scores. Notably, those scoring zero on the CRT (i.e., those presumed to be most reliant on intuition) apparently did not benefit from the hint at all.

DISCUSSION

The results of Study 2 support the dualsystems hypothesis and suggest that one reason people fail to maximize is that they fail to consider the maximizing strategy as an alternative. In other words, probability matching is so intuitive that it is often the only solution to the problem that participants consider. When the maximizing strategy was provided in a hint prior to the choice task, participants not only endorsed the strategy but they also went on to use the strategy more than those who did not receive the hint prior to the task. This suggests that, for many participants, the problem is not in their ability to realize that maximizing is superior, but in their likelihood of considering maximizing in the first place.

This finding answers our initial question regarding why choice-based studies and strategy-based studies differ, by suggesting that it is because the latter provides both strategies for participants to evaluate. In other words, providing a hint avoids the discrepancy in strategy availability typical of choice-based studies.

The discrepancy in strategy availability suggests that matching is intuitive whereas maximizing requires deliberation. Specifically, the deliberative system must override the impulse to employ the first strategy that comes to mind and engage in a more exhaustive search for alternative strategies. The importance of deliberation is further supported by the correlation between CRT scores and maximizing observed in Study 2. A correlation that was maintained even when controlling for ability at mathematics.These findings suggested that those who were more deliberative in nature were more likely to maximize.

It is not clear how any of these findings would be easily explained by the patternsearch hypothesis. This account does not answer (a) why participants would be less likely to search for patterns (and thus maximize) after being provided with both the maximizing and matching strategies in a hint or (b) why those with a tendency to deliberate (as indicated by CRT score) would not bother searching for patterns. Thus, the results of Study 2 provide more support for a dualsystems account than for a pattern search account of probability matching.

Dualsystems theory has provided a compelling account of the data presented thus far, but one could argue that it misses the real crux of the issue by failing to address why probability matching is intuitive in the first place. It also fails to explain why groups (non-human species and young children) that we generally consider to be less deliberative than adult humans choose maximizing strategies.

To address these issues we have suggested a new theory that we call expectation matching.Expectation matching is based on research showing that people tend to expect very small samples to be representative of the contingencies of their parent populations (Tversky&Kahenman 1971; Kahenman&Tversky 1972). For example, if you were to flip a coin twice you would expect it to come up heads once and tails once, although the chances that it will come up as one or the other twice in a row are quite high. The representativeness bias is the tendency to overestimate the likelihood that the outcome will consist of one heads and one tails. In the die problem mentioned above, this means that participants would expect 7 green and 3 red rolls exactly, even when you only roll the die 10 times. Expectation matching argues that the probability matching strategy is selected based on this highly salient expected outcome.

Work by Gal and Barron (1996) provides support for the hypothesis that this expectation is highly salient to participants employing a probability matching strategy. They asked participants to explain why they selected certain strategies over others after participants completed a word problem version of a binary choice task. Participants that chose to probability match tended to citeas the reasons for their strategy choice(1) expectations of the long term distribution, and (2) the belief that the distribution will turn up as predicted.

Gal and Barron(1996) also noted that these participants tended to have a poorer understanding of independence, mostly due to their tendency to endorse the gambler's fallacy[5]. Gambler's fallacy can also be attributed to the representativeness bias. Those that endorse the gambler's fallacy tend to view runs as increasing the likelihood that the non-run event will occur "on the grounds that the sequence will balance out locally" (participant's words in Gal& Barron,

[5]This is in spite of the fact that participants were usually able to define independence, although they often didnot trust their own mathematical intuitions and went with their "gut."

1996, p. 88). These endorsements of the Gambler's fallacy led many participants to select a strategy in which they mostly guessed the more probable outcome, but switched to the less probable outcome after an unusually long streak of the more probable occurrence. This resulted in a number of participants classified as over-matchers.[6] The focus of both over-matchers and probability matchers on the local representativeness of the sequence suggests that these participants are heavily focused on global expected outcome (overall, how many reds and greens do I expect) rather than on the contingencies of each individual die roll.

With the expectation of 7 green and 3 red beinghighly salient, subjects then match to this expectation. But why? One possibility is that they fail to properly identify the actual problem that they are attempting to solve. To predict outcomes perfectly in a dual choice task, participants do not need to know the overall distribution of green and red (the expectation), but instead the *sequence* of these reds and greens. It is this second part of the problem that participants overlook.Thus, they substitute for the answer to a relatively easy question (Which overall distribution is most likely?) the answer to the more difficult question (Which exact sequence of reds and greens is most likely?), a process known as attribute substitution (Kahneman& Frederick, 2002). Of course, the answer to the easy question is 7 green and 3 red, but to actually accurately predict outcomes in a binary choice task, one needs to know the answer to the harder question. In this case, the most likely single sequence is all green[7], because green is the most likely occurrence on any individual roll. Adding a red to the sequence adds a lower probability event, decreasing the overall probability of that particular sequence. But if you

[6] Note that this might explain the earlier findings of Wolford et al.(2004), in which making a sequence with more alterations (and therefore fewer streaks) led to significantly more maximizing, due to over-matchers becoming maximizers.

[7] All green should occur close to 3% of the time, while any give sequence of 7 green and 3 red should only occur roughly 0.2% of the time.

ignore sequence information and simply ask the easy question, it is easy to think that you need only match what you expect to see overall to achieve the highest payoff.

The expectation matching hypothesis also provides a solution to the logical inconsistency of why a more deliberative behavior (maximizing) is endorsed by young children and non-human mammals, while the intuitive behavior (matching) is not. This puzzle is not easily solved by the dual systems account alone, nor is it addressed by the patternsearch hypothesis. Expectation matching, however, would claim that young children and non-human mammals are not able to generate an expectation across multiple trials, thus they should not be subject to the representativeness bias. According to this hypothesis, a young child, for example, would maximize in the die game, but would not be able to report that they expect to see 7 green and 3 red. Congruent with this prediction, DerksandPaclisanu (1967) found that while 4-year-olds did maximize, they were not able to verbally report which outcome was more likely. The most likely explanation for maximizing behavior in young children and non-human mammals, therefore, would be that they choose the more likely outcome all the time because it is the most frequently rewarded behavior. If they are able to generate some kind of expectation, without the ability to plan (and recall) across multiple trials, they would likely ask themselves "Which is more likely on *just this* trial?" The answer would always be green and, therefore, their choices would always be green.

Studies 3, 4 and 5are inspired by this localized view. We endeavored to test the expectation matching hypothesis by putting undergraduate students in the mind frame of a rat. In other words, we wanted to reduce their tendency to generate an expectation for the entire sequence of 10 outcomes, by focusing them on individual trials instead of the whole sequence. Expectation matching would predict that this should increase maximizing behavior.

We also examined whether the hard question “which sequence is most likely?” posed above is actually hard. We probed whether participants are able to identify that all green is the single most likely sequence by asking participants “Which strategy (matching or maximizing) is most likely to get all 10 guesses correct?”

Study 3

METHOD

Participants

Participants were 132 undergraduates from the University of Waterloo who completed the study online for course credit.

Procedure

They were asked to play a series of 10 guessing games in which they were to predict for each game which of two possible outcomes would occur. They were told in advance that, in each of the 10 games, one outcome would occur with 70% probability and the other with 30% probability. Participants were asked to imagine that they would be paid $1 for each correct prediction, and asked to indicate their prediction in each of the 10 games. They were not informed of the game outcomes.

In the unique games condition, participants were told they would be playing 10 *different* games, the specifics of which would be described to them as they played each game. They then made predictions for 10 mathematically equivalent games with superficial individuating characteristics. For example, one game involved drawing ping-pong balls from a bingo cage; another involved spinning a wheel of fortune; and another involved rolling a 10-sided die. Order of game presentation was randomized across participants. In the repeated games condition, one of the 10 games from the unique games condition was randomly selected for each participant and presented to him or her 10 times.

Finally, participants answered a strategy question in which they compared the maximizing and matching strategies. As in Studies 1 and 2, they were asked which of the two strategies was most likely to earn the most money. Now, however, we also asked them which of the two strategies was most likely to result in getting all 10 guesses correct.

RESULTS

Data from two participants who predicted the unlikely outcome more often than the likely outcome were excluded from further analysis, as were data from three additional participants who did not complete the entire choice task.

Figure 3 shows the full distribution of the number of times that each participant predicted the more likely outcome. Participants in the unique games condition (n = 66) were more likely than those in the repeated games condition (n = 61) to engage in strict maximization, that is, to predict the more likely outcome in all ten choices, χ^2 (1, n = 127) = 8.39*p* = .004. Strict probability matching (predicting the more likely outcome on 7 of 10 choices), by contrast, was common in the repeated games condition but not in the unique games condition, χ^2 (1, n = 127) = 24.11*p*< .001.

When asked which strategy was more likely to earn the most money, 86.6% of participants in the repeated condition answered maximizing. By contrast, only 45.9% of those same participants thought maximizing was the best way to get all ten guesses correct. The numbers were higher in the unique condition, with 92.4% responding that maximizing was the best way to make money and 71.2% responding that it was the best way to get all ten guesses correct. A repeated measures ANOVA revealed that there was a significant main effect of question, $F(1,125)=33.33$, $MSE=0.17$, $p<.001$, and a significant main effect of condition $F(1,125)=11.30$, $MSE=0.196$, $p<.001$, but no interaction.

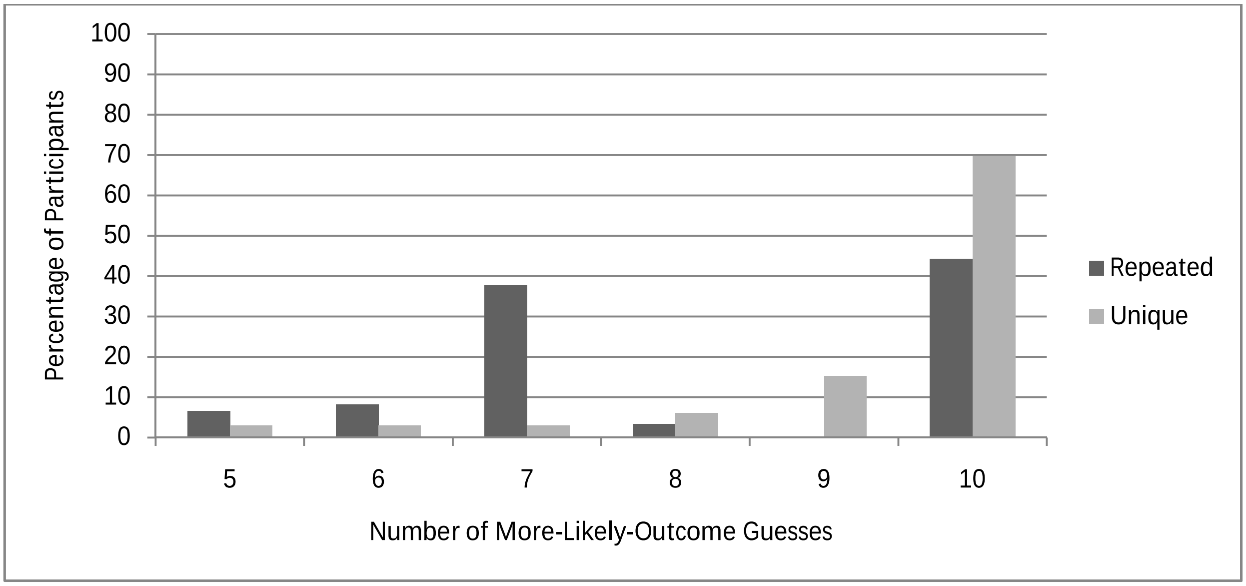

Fig. 3.*Distribution of the number of predictions in study 3 of the morelikely outcome (rather than the lesslikely outcome) across participants in the two conditions.*

Study 4

In Study 3, the various games presented in the unique games condition all shared common outcome probabilities of 70% and 30%, and this was noted in the instructions, but how these probabilities were presented varied across the games. In some, they were given directly as probabilities or relative frequencies of occurrence; in others, they could be inferred from characteristics of the chance device (e.g., a draw from a bag of 100 tickets, where 70 are orange and the other 30 are black). Study 4 tested whether individuating the outcomes in the sequence would decrease probability matching and increase maximizing even when the games wereall based on a common chance device.

METHOD

Participants

University of Waterloo undergraduates (n= 129) completed the online study in exchange for psychology course credit.

Procedure

The participants were asked to consider a game involving 10 rolls of a 10-sided die. Seven sides of the die, they were told, were marked one way, and the remaining three sides were marked another way. In the repeated games condition, all 10 rolls involved a die with 7 red sides and 3 green sides. In the unique games condition, each roll was said to involve a different die with unique markings. For instance, in addition to the die marked with red and green sides as above, another die had 7 sides marked with triangles and 3 sides marked with squares. Other markings included letters and icons such as flowers and hearts.

Participants responded to the same two strategy questions as mentioned in Study 3.

RESULTS

Data from one participant who predicted the unlikely outcome more often than the likely outcome were excluded from further analysis, as were data from a second participant who did not complete the entire choice task.

Within the repeated games condition, for some participants the sequence was described as a single game consisting of 10 rolls of the die, and for others it was described as 10 games of 1 roll each. Analyses demonstrated that this factor had no influence on the results and so it is not discussed further.

Figure 4 shows the distribution of the number of times each participant guessed the more likely outcome. Participants in the unique games condition (n = 38) were more likely than those in the repeated games condition (n = 89) to engage in strict maximization, that is, to predict the more likely outcome in all ten choices, χ^2 (1, n= 127) = 4.59, $p = .03$. Strict probability matching, by contrast, was more prevalent in the repeated games condition than in the unique games condition, χ^2 (1, n = 127) = 5.41, $p = .02$.

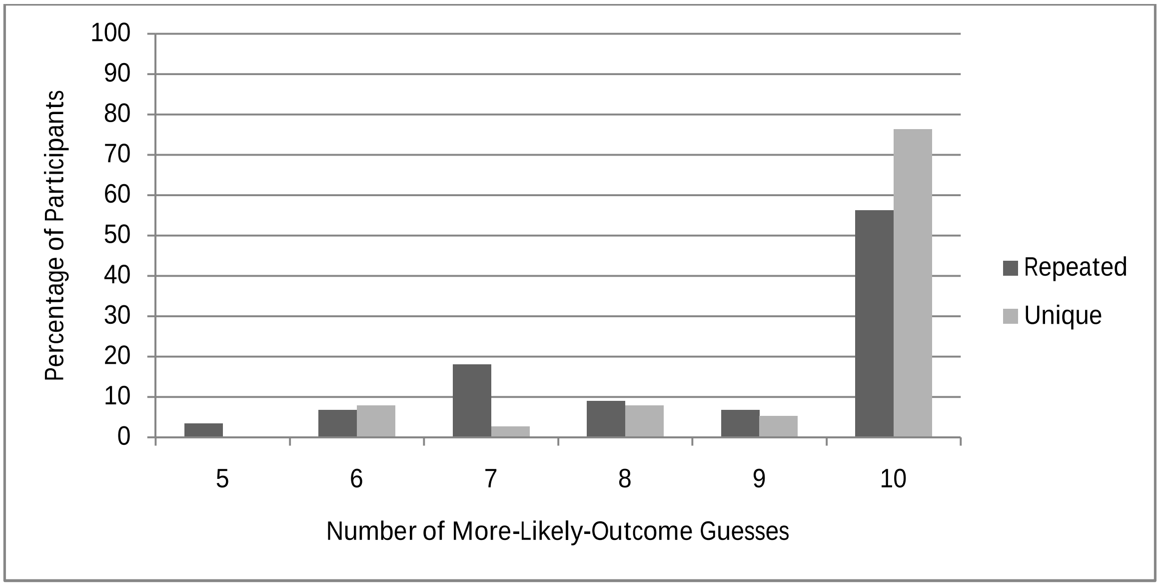

Fig. 4.*Distribution of the number of predictions of the morelikely outcome (rather than the lesslikely outcome) across participants in the two conditions.*

Once again, the strategy questions revealed that participants were more able to identify maximizing as a better strategy when asked about making money than when asked about getting all 10 guesses correct. In the repeated condition, 85.4% of participants responded that maximizing was the best strategy for making money whereas only 60.7% also chose it as the best strategy for getting all 10 guesses correct. In the unique condition, the numbers rose to 94.7% and 65.8%, respectively. A repeated measures ANOVA revealed a main effect of question $F(1,125)=28.57$, $MSE=0.113$, $p<.000$. There was no main effect of condition, nor was there an interaction.

Study 5

In Studies 3 and 4, introducing individuating features to the outcome sequence reduced probability matching and increased maximizing. We suggest that this is because the individuating features made it less likely that participants would generate and expectation for all 10 guesses at once. In Study 5, we sought convergent evidence for this interpretation by keeping the features of the sequence itself fixed but preceding it with a priming manipulation designed to draw focus either to the sequence as a whole or to the individual outcomes within the sequence.

METHOD

Participants

University of Waterloo students (N = 84) were recruited at a public location on campus and asked to complete a brief questionnaire.

Procedure

Participants were asked to consider the game from Study 4 in which a 10-sided die, with 7 green sides and 3 red sides, would be rolled 10 times and each correct prediction would yield a $1 payoff. Participants indicated their predictions for each of the 10 rolls by circling either green or redon a grid.

Before playing the game, those in the global focus condition were asked, “In 10 rolls of the die, how many times would you expect each outcome?” This question was intended to encourage generation of an aggregate expectation for the sequence as a whole. Participants in the local focus condition were asked, “On any individual roll of the die, which color is more likely to be rolled?” This question was intended to focus participants on the contingencies of an individual trial rather than expected outcomes for the sequence as a whole, and thereby to reduce the tendency to probability match. Over 80% of participants in the global focus condition and all participants in the local focus condition answered the priming question correctly.

RESULTS

Incomplete data from one participant were removed from further analysis.

Figure 5 shows the distribution of the number of times each participant guessed the more likely outcome. Participants in the local focus condition (n = 42) were more likely than those in the global focus condition (n = 41) to engage in strict maximization, that is, to predict the more likely outcome in all ten choices, χ^2 (1, n= 83) = 7.33, p = .007. Strict probability matching, by contrast, was more prevalent in the global focus condition than in the local focus condition, χ^2 (1, n= 83) = 4.35, p = .04.

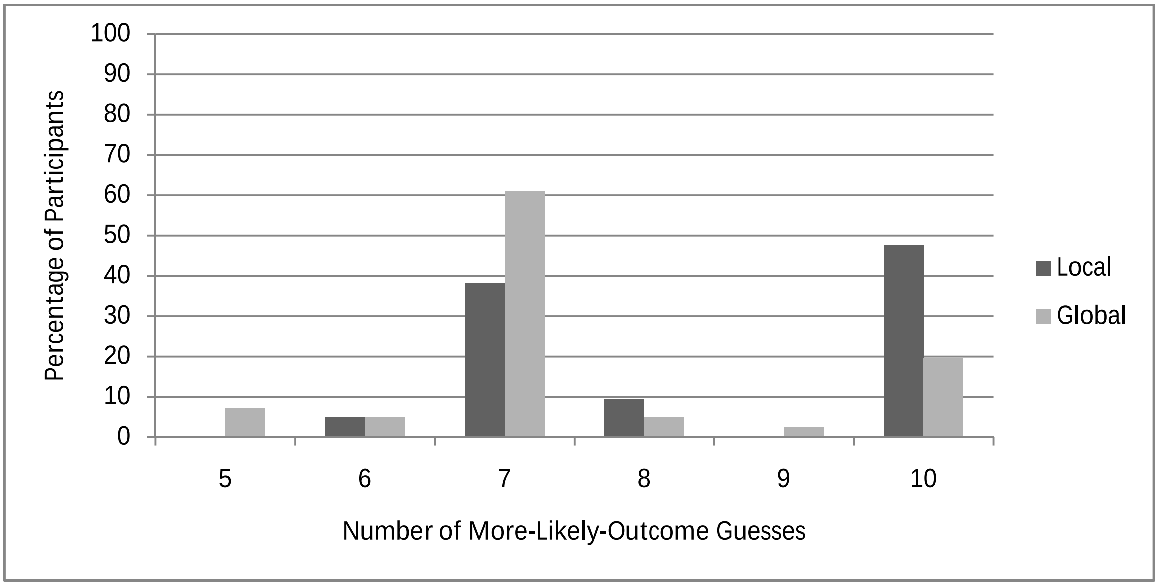

Fig. 5.*Distribution of the number of predictions in study 5 of the morelikely outcome (rather than the lesslikely outcome) across participants in the two conditions.*

Discussion of Studies 3, 4 and 5

Results from Studies 3 and 4 showed that making trials of a binary choice task appear to be independent by (a) taking each trial from a superficially different game or (b) rolling a superficially different die for each roll leads to more maximizing behavior. Furthermore, Study 5 suggests that even focusing participants on individual trials as opposed to global outcomes can increase the rate of maximization. In addition, Studies 3 and 4 demonstrated that participants have much more difficulty realizing that maximizing is the best strategy for predicting all 10 guesses correctly than they do acknowledging it as the most likely strategy to win the most money. This is especially amusing because, in this particular case, these two questions are based on the same underlying concept. The strategy that is most likely to make you the most money is also, by definition, the most likely to get all guesses correct.

These results are congruent with those predicted by the expectation matching hypothesis, but should be viewed as preliminary evidence. It is important to note that the expectation matching hypothesis makes a particular prediction about the scenarios tested in these studies, and that those predictions were confirmed, but the studies did not directly test whether expectations were affected by the manipulation. Much further research is required to confirm the exact effect of these manipulations

I have four specific recommendations for further research. First, these studies do not establish that individuating trials actually operates by affecting expectations or reducing representativeness bias. Mechanism was not addressed in these studies so further research is necessary to establish exactly what mechanism is operating. Preliminary work examining this issue could simply ask participants some probing questions about why they chose particular strategies, similar to what was done by Gal and Barron (1996), but in the context of individuated trials. Those in the individuated conditions should make less reference to the expected outcome across all 10 trials and should be less susceptible to fallacies based on a poor understanding of independence, such as gambler's fallacy.

Second, even if we accept that expectation generation is affected by individuating the trials, this work still does not clarify whether individuating trials prevents participants from generating an expectation in the first place or whether it simply encourages them to ignore the expectation. One possible way to resolve this issue would be to cross the unique repeated manipulations from Studies 3 and 4 with the priming manipulation(plus a control condition with no prime) from Study 5. The participants of interest in this design would be those who were in the unique condition and received the global prime. If the unique manipulation prevents the generation of an expectation, than forcing them to generate an expectation via the global prime should increased matching in that condition relative to a control condition. By contrast, if the expectation was always generated but simply ignored, the global prime should have no effect on strategy selection.

Third, in Study 5 we did not include a control condition because we were interested simply in comparing those with a local focus to those with a global focus. However, if we are interested in establishing that a global focus is the default, it would be useful to rerun Study 5 with a control condition that receives no priming question. These subjects should look very similar to those in the global condition.

Finally, the strategy questions in Studies 3 and 4 provided a preliminary investigation of participants' understanding of the proposed "hard question," namely "Which exact sequence is most likely?" We demonstrated that answers to this question are much poorer than to the theoretically identical question querying making money—and indeed also worse than maximizing rates in the game itself. It would be useful to have further research confirming that understanding of the hard question is poor. Furthermore, future research should address why answers to the sequence question are worse than answers to the money question even though, in this case, they are asking the same thing as the money question.

Conclusion

Overall, this thesis found that thepatternsearch hypothesis does not provide a compelling account of the results obtained in Studies 1-5. Study 1 demonstrated that probability matching behavior is not influenced by the availability of pattern information, nor is it affected by a lack of ability to include pattern information in responses, contrary to what would be predicted by the pattern-search account. Study 2 found that making both strategies equally available greatly improved participants' ability to choose the optimal strategy. This suggests that probability matching may occur because participants fail to think of maximizing as an alternative. Again this is not easily accounted for by the patternsearch account, but fits well with the dualsystems account of probability matching. This difference in availability also explains why studies in the literature that make both strategies equally available find higher rates of maximizing.

The dualsystems account of probability matching argues that maximizing is a deliberative strategy that must override the more intuitive probability matching strategy, but does not explain why probability matching is intuitive. I suggest that this intuition arises from expectation matching, the tendency to generate an expectation about aggregate outcomes across a sequence and to make choices based on this expectation. I hypothesized that this expectation influences choice through attribute substitution, with participants focusing on the most probable aggregate outcome and ignoring the fact that they lack information about how to sequence their guesses.

Studies 3, 4, and 5 demonstrated that making trials more independent greatly reduces probability matching, as predicted by the expectation matching hypothesis. In addition, Studies 3 and 4 also established that participants are very poor at answering the relatively difficult question regarding outcome sequence. This supports the argument that participants tend to ignore this information and substitute the easier question of expected outcomes, as suggested by attribute substitution.

Further research is required to test whether expectation matching is a good extension of the dualsystems account of probability matching. Nevertheless, the results of the studies reported here strongly suggest that dual systems is a better account of probability matching behavior than is pattern matching, that probability matching is intuitive, and that expectation matching may be the mechanism behind the intuition.

Letters of copyright permission

ELSEVIER LICENSE
TERMS AND CONDITIONS

Jan 09, 2012

Supplier	Elsevier Limited The Boulevard,Langford Lane Kidlington,Oxford,OX5 1GB,UK
Registered Company Number	1982084
Customer name	Greta James
Customer address	404 Lakeview Dr.
	Waterloo, ON N2L4Z6
License number	2824970402294
License date	Jan 09, 2012
Licensed content publisher	Elsevier
Licensed content publication	Cognition
Licensed content title	Probability matching in choice under uncertainty: Intuition versus deliberation
Licensed content author	Derek J. Koehler,Greta James
Licensed content date	October 2009
Licensed content volume number	113
Licensed content issue number	1
Number of pages	5
Start Page	123
End Page	127
Type of Use	reuse in a thesis/dissertation
Intended publisher of new work	other
Portion	excerpt

Number of excerpts	1
Format	both print and electronic
Are you the author of this Elsevier article?	Yes
Will you be translating?	No
Order reference number	
Title of your thesis/dissertation	Betting on the Unexpected: The Effect of Expectation Matching on Choice Strategies in a Binary Choice Task
Expected completion date	Feb 2012
Estimated size (number of pages)	49
Elsevier VAT number	GB 494 6272 12
Permissions price	0.00 USD
VAT/Local Sales Tax	0.0 USD / 0.0 GBP
Total	0.00 USD
Terms and Conditions	

INTRODUCTION

1. The publisher for this copyrighted material is Elsevier. By clicking "accept" in connection with completing this licensing transaction, you agree that the following terms and conditions apply to this transaction (along with the Billing and Payment terms and conditions established by Copyright Clearance Center, Inc. ("CCC"), at the time that you opened your Rightslink account and that are available at any time at http://myaccount.copyright.com).

GENERAL TERMS

2. Elsevier hereby grants you permission to reproduce the aforementioned material subject to the terms and conditions indicated.

3. Acknowledgement: If any part of the material to be used (for example, figures) has appeared in our publication with credit or acknowledgement to another source, permission must also be sought from that source. If such permission is not obtained then that material may not be included in your publication/copies. Suitable acknowledgement to the source must be made, either as a footnote or in a reference list at the end of your publication, as follows:

“Reprinted from Publication title, Vol /edition number, Author(s), Title of article / title of chapter, Pages No., Copyright (Year), with permission from Elsevier [OR APPLICABLE SOCIETY COPYRIGHT OWNER].” Also Lancet special credit - “Reprinted from The Lancet, Vol. number, Author(s), Title of article, Pages No., Copyright (Year), with permission from Elsevier.”

4. Reproduction of this material is confined to the purpose and/or media for which permission is hereby given.

5. Altering/Modifying Material: Not Permitted. However figures and illustrations may be altered/adapted minimally to serve your work. Any other abbreviations, additions, deletions and/or any other alterations shall be made only with prior written authorization of Elsevier Ltd. (Please contact Elsevier at permissions@elsevier.com)

6. If the permission fee for the requested use of our material is waived in this instance, please be advised that your future requests for Elsevier materials may attract a fee.

7. Reservation of Rights: Publisher reserves all rights not specifically granted in the combination of (i) the license details provided by you and accepted in the course of this licensing transaction, (ii) these terms and conditions and (iii) CCC's Billing and Payment terms and conditions.

8. License Contingent Upon Payment: While you may exercise the rights licensed immediately upon issuance of the license at the end of the licensing process for the transaction, provided that you have disclosed complete and accurate details of your proposed use, no license is finally effective unless and until full payment is received from you (either by publisher or by CCC) as provided in CCC's Billing and Payment terms and conditions. If full payment is not received on a timely basis, then any license preliminarily granted shall be deemed automatically revoked and shall be void as if never granted. Further, in the event that you breach any of these terms and conditions or any of CCC's Billing and Payment terms and conditions, the license is automatically revoked and shall be void as if never granted. Use of materials as described in a revoked license, as well as any use of the materials beyond the scope of an unrevoked license, may constitute copyright infringement and publisher reserves the right to take any and all action to protect its copyright in the materials.

9. Warranties: Publisher makes no representations or warranties with respect to the licensed material.

10. Indemnity: You hereby indemnify and agree to hold harmless publisher and CCC, and their respective officers, directors, employees and agents, from and against any and all claims arising out of your use of the licensed material other than as specifically authorized pursuant to this license.

11. No Transfer of License: This license is personal to you and may not be sublicensed, assigned, or transferred by you to any other person without publisher's written permission.

12. No Amendment Except in Writing: This license may not be amended except in a writing signed by both parties (or, in the case of publisher, by CCC on publisher's behalf).

13. Objection to Contrary Terms: Publisher hereby objects to any terms contained in any purchase order, acknowledgment, check endorsement or other writing prepared by you,

which terms are inconsistent with these terms and conditions or CCC's Billing and Payment terms and conditions. These terms and conditions, together with CCC's Billing and Payment terms and conditions (which are incorporated herein), comprise the entire agreement between you and publisher (and CCC) concerning this licensing transaction. In the event of any conflict between your obligations established by these terms and conditions and those established by CCC's Billing and Payment terms and conditions, these terms and conditions shall control.

14. Revocation: Elsevier or Copyright Clearance Center may deny the permissions described in this License at their sole discretion, for any reason or no reason, with a full refund payable to you. Notice of such denial will be made using the contact information provided by you. Failure to receive such notice will not alter or invalidate the denial. In no event will Elsevier or Copyright Clearance Center be responsible or liable for any costs, expenses or damage incurred by you as a result of a denial of your permission request, other than a refund of the amount(s) paid by you to Elsevier and/or Copyright Clearance Center for denied permissions.

LIMITED LICENSE

The following terms and conditions apply only to specific license types:

15. **Translation**: This permission is granted for non-exclusive world **English** rights only unless your license was granted for translation rights. If you licensed translation rights you may only translate this content into the languages you requested. A professional translator must perform all translations and reproduce the content word for word preserving the integrity of the article. If this license is to re-use 1 or 2 figures then permission is granted for non-exclusive world rights in all languages.

16. **Website**: The following terms and conditions apply to electronic reserve and author websites:
Electronic reserve: If licensed material is to be posted to website, the web site is to be password-protected and made available only to bona fide students registered on a relevant course if:
This license was made in connection with a course,
This permission is granted for 1 year only. You may obtain a license for future website posting,
All content posted to the web site must maintain the copyright information line on the bottom of each image,
A hyper-text must be included to the Homepage of the journal from which you are licensing athttp://www.sciencedirect.com/science/journal/xxxxx or the Elsevier homepage for books athttp://www.elsevier.com , and
Central Storage: This license does not include permission for a scanned version of the material to be stored in a central repository such as that provided by Heron/XanEdu.

17. **Author website** for journals with the following additional clauses:

All content posted to the web site must maintain the copyright information line on the bottom of each image, and
he permission granted is limited to the personal version of your paper. You are not allowed to download and post the published electronic version of your article (whether PDF or HTML, proof or final version), nor may you scan the printed edition to create an electronic version,
A hyper-text must be included to the Homepage of the journal from which you are licensing athttp://www.sciencedirect.com/science/journal/xxxxx , As part of our normal production process, you will receive an e-mail notice when your article appears on Elsevier's online service ScienceDirect (www.sciencedirect.com). That e-mail will include the article's Digital Object Identifier (DOI). This number provides the electronic link to the published article and should be included in the posting of your personal version. We ask that you wait until you receive this e-mail and have the DOI to do any posting.
Central Storage: This license does not include permission for a scanned version of the material to be stored in a central repository such as that provided by Heron/XanEdu.

18. **Author website** for books with the following additional clauses:
Authors are permitted to place a brief summary of their work online only.
A hyper-text must be included to the Elsevier homepage at http://www.elsevier.com

All content posted to the web site must maintain the copyright information line on the bottom of each image
You are not allowed to download and post the published electronic version of your chapter, nor may you scan the printed edition to create an electronic version.
Central Storage: This license does not include permission for a scanned version of the material to be stored in a central repository such as that provided by Heron/XanEdu.

19. **Website** (regular and for author): A hyper-text must be included to the Homepage of the journal from which you are licensing
at http://www.sciencedirect.com/science/journal/xxxxx. or for books to the Elsevier homepage at http://www.elsevier.com

20. **Thesis/Dissertation**: If your license is for use in a thesis/dissertation your thesis may be submitted to your institution in either print or electronic form. Should your thesis be published commercially, please reapply for permission. These requirements include permission for the Library and Archives of Canada to supply single copies, on demand, of the complete thesis and include permission for UMI to supply single copies, on demand, of the complete thesis. Should your thesis be published commercially, please reapply for permission.

21. **Other Conditions**:

v1.6

If you would like to pay for this license now, please remit this license along with your payment made payable to "COPYRIGHT CLEARANCE CENTER" otherwise you will be invoiced within 48 hours of the license date. Payment should be in the form of a check or money order referencing your account number and this invoice number RLNK500695574.
Once you receive your invoice for this order, you may pay your invoice by credit card. Please follow instructions provided at that time.

Make Payment To:
Copyright Clearance Center
Dept 001
P.O. Box 843006
Boston, MA 02284-3006

For suggestions or comments regarding this order, contact RightsLink Customer Support:customercare@copyright.com or +1-877-622-5543 (toll free in the US) or +1-978-646-2777.

Gratis licenses (referencing $0 in the Total field) are free. Please retain this printable license for your reference. No payment is required.

SPRINGER LICENSE
TERMS AND CONDITIONS

Jan 09, 2012

This is a License Agreement between Greta James ("You") and Springer ("Springer") provided by Copyright Clearance Center ("CCC"). The license consists of your order details, the terms and conditions provided by Springer, and the payment terms and conditions.

All payments must be made in full to CCC. For payment instructions, please see information listed at the bottom of this form.

License Number

2824960510241

License date

Jan 09, 2012

Licensed content publisher

Springer

Licensed content publication

Memory & Cognition

Licensed content title

Probability matching and strategy availability

Licensed content author

Derek J. Koehler

Licensed content date

Jan 1, 2010

Volume number

38

Issue number

6

Type of Use

Thesis/Dissertation

Portion

Excerpts

Author of this Springer article

Yes and you are a contributor of the new work

Order reference number

Title of your thesis / dissertation

Betting on the Unexpected: The Effect of Expectation Matching on Choice Strategies in a Binary Choice Task

Expected completion date

Feb 2012

Estimated size(pages)

49

Total

0.00 CAD

Terms and Conditions

Introduction
The publisher for this copyrighted material is Springer Science + Business Media. By clicking "accept" in connection with completing this licensing transaction, you agree that the following terms and conditions apply to this transaction (along with the Billing and Payment terms and conditions established by Copyright Clearance Center, Inc. ("CCC"), at the time that you opened your Rightslink account and that are available at any time at http://myaccount.copyright.com).

Limited License
With reference to your request to reprint in your thesis material on which Springer Science and Business Media control the copyright, permission is granted, free of charge, for the use indicated in your enquiry. Licenses are for one-time use only with a maximum distribution equal to the number that you identified in the licensing process.

This License includes use in an electronic form, provided it is password protected or on the university's intranet, destined to microfilming by UMI and University repository. For any other electronic use, please contact Springer at (permissions.dordrecht@springer.com or permissions.heidelberg@springer.com)

The material can only be used for the purpose of defending your thesis, and with a maximum of 100 extra copies in paper.

Although Springer holds copyright to the material and is entitled to negotiate on rights, this license is only valid, provided permission is also obtained from the (co) author (address is given with the article/chapter) and provided it concerns original material which does not carry references to other sources (if material in question appears with credit to another source, authorization from that source is required as well). Permission free of charge on this occasion does not prejudice any rights we might have to charge for reproduction of our copyrighted material in the future.

Altering/Modifying Material: Not Permitted
However figures and illustrations may be altered minimally to serve your work. Any other abbreviations, additions, deletions and/or any other alterations shall be made only with prior written authorization of the author(s) and/or Springer Science + Business Media. (Please contact Springer at permissions.dordrecht@springer.com or permissions.heidelberg@springer.com)

Reservation of Rights
Springer Science + Business Media reserves all rights not specifically granted in the combination of (i) the license details provided by you and accepted in the course of this licensing transaction, (ii) these terms and conditions and (iii) CCC's Billing and Payment terms and conditions.

Copyright Notice:
Please include the following copyright citation referencing the publication in which the material was originally published. Where wording is within brackets, please include verbatim.
"With kind permission from Springer Science+Business Media: <book/journal title, chapter/article title, volume, year of publication, page, name(s) of author(s), figure number(s), and any original (first) copyright notice displayed with material>."

Warranties: Springer Science + Business Media makes no representations or warranties with respect to the licensed material.

Indemnity
You hereby indemnify and agree to hold harmless Springer Science + Business Media and CCC, and their respective officers, directors, employees and agents, from and against any and all claims arising out of your use of the licensed material other than as specifically authorized pursuant to this license.

No Transfer of License
This license is personal to you and may not be sublicensed, assigned, or transferred by you to any other person without Springer Science + Business Media's written permission.

No Amendment Except in Writing
This license may not be amended except in a writing signed by both parties (or, in the case of Springer Science + Business Media, by CCC on Springer Science + Business Media's behalf).

Objection to Contrary Terms
Springer Science + Business Media hereby objects to any terms contained in any purchase order, acknowledgment, check endorsement or other writing prepared by you, which terms are inconsistent with these terms and conditions or CCC's Billing and Payment terms and conditions. These terms and conditions, together with CCC's Billing and Payment terms and conditions (which are incorporated herein), comprise the entire agreement between you and Springer Science + Business Media (and CCC) concerning this licensing transaction. In the event of any conflict between your obligations established by these terms and conditions and those established by CCC's Billing and Payment terms and conditions, these terms and conditions shall control.

Jurisdiction
All disputes that may arise in connection with this present License, or the breach thereof, shall be settled exclusively by the country's law in which the work was originally published.

Other terms and conditions:

v1.2

If you would like to pay for this license now, please remit this license along with your payment made payable to "COPYRIGHT CLEARANCE CENTER" otherwise you will be invoiced within 48 hours of the license date. Payment should be in the form of a check or money order referencing your account number and this invoice number RLNK500695549.
Once you receive your invoice for this order, you may pay your invoice by credit card. Please follow instructions provided at that time.

Make Payment To:
Copyright Clearance Center
Dept 001
P.O. Box 843006
Boston, MA 02284-3006

For suggestions or comments regarding this order, contact RightsLink Customer Support:customercare@copyright.com or +1-877-622-5543 (toll free in the US) or +1-978-646-2777.

Gratis licenses (referencing $0 in the Total field) are free. Please retain this printable license for your reference. No payment is required.

References

Brackbill, Y., & Bravos, A. (1962). Supplementary report: The utility of correctly predicting infrequent events. *Journal of Experimental Psychology, 64*(6), 648-649. doi:10.1037/h0046489

Brackbill, Y., Kappy, M. S., & Starr, R. H. (1962). Magnitude of reward and probability learning. *Journal of Experimental Psychology, 63*(1), 32-35. doi:10.1037/h0045051

Bullock, D. H., & Bitterman, M. E. (1962). Probability-matching in the pigeon. *The American Journal of Psychology, 75*(4), 634-634-639. doi:10.2307/1420288

Derks, P. L., & Paclisanu, M. I. (1967). Simple strategies in binary prediction by children and adults. *Journal of Experimental Psychology, 73*(2), 278-285. doi:10.1037/h0024137

Edwards, W. (1961). Probability learning in 1000 trials. *Journal of Experimental Psychology, 62*(4), 385-394. doi:10.1037/h0041970

Falk, R., & Konold, C. (1997). Making sense of randomness: Implicit encoding as a basis for judgment. *Psychological Review, 104*(2), 301-301-318. doi:10.1037/0033-295X.104.2.301

Fantino, E., & Esfandiari, A. (2002). Probability matching: Encouraging optimal responding in humans. *Canadian Journal of Experimental Psychology/Revue Canadienne De Psychologie Expérimentale, 56*(1), 58-63. doi:10.1037/h0087385

Frederick, S. (2005). Cognitive reflection and decision making. *Journal of Economic Perspectives, 19*, 24-42.

Gal, I., & Baron, J. (1996). Understanding repeated simple choices. *Thinking & Reasoning, 2*(1), 81-98. doi:10.1080/135467896394573

Goodnow, J. J. (1955). Determinants of choice-distribution in two-choice situations. *The American Journal of Psychology, 68*, 106-116. doi:10.2307/1418393

Goodnow, J. J. (1955). Response sequences in a pair of two-choice probability situations. *American Journal of Psychology, 68*, 624-630. doi:10.2307/1418790

Grant, D. A., Hake, H. W., & Hornseth, J. P. (1951). Acquisition and extinction of a verbal conditioned response with differing percentages of reinforcement. *Journal of Experimental Psychology, 42*(1), 1-1-5. doi:10.1037/h0054051

Kahneman, D., & Frederick, S. (2002). Representativeness revisited: Attribute substitution in intuitive judgment. In T. Gilovich, D. Griffin & D. Kahneman (Eds.), (pp. 49-49-81). New York, NY, US: Cambridge University Press.

Kahneman, D., & Frederick, S. (2005). In Holyoak K. J., Morrison R. G. (Eds.), *A model of heuristic judgment*. New York, NY, US: Cambridge University Press.

Kahneman, D., & Tversky, A. (1972). Subjective probability: A judgment of representativeness. *Cognitive Psychology, 3*(3), 430-454. doi:10.1016/0010-0285(72)90016-3

Kahneman, D., & Tversky, A. (1979). On the interpretation of intuitive probability: A reply to jonathan cohen. *Cognition, 7*(4), 409-409-411. doi:10.1016/0010-0277(79)90024-6

Kogler, C., & Kuhberger, A. (2007). Dual process theories: A key for understanding diversification bias? *Journal of Risk and Uncertainty, 34*, 145-154.

Lee, D., Conroy, M. L., McGreevy, B. P., & Barraclough, D. J. (2004). Reinforcement learning and decision making in monkeys during a competitive game. *Cognitive Brain Research, 22*(1), 45-58. doi:10.1016/j.cogbrainres.2004.07.007

Lopes, L. L., & Oden, G. C. (1987). Distinguishing between random and nonrandom events. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 13*(3), 392-392-400. doi:10.1037/0278-7393.13.3.392

Miller, M. B., Valsangkar-Smyth, M., Newman, S., Dumont, H., & Wolford, G. (2005). Brain activations associated with probability matching. *Neuropsychologia, 43*(11), 1598-1608. doi:10.1016/j.neuropsychologia.2005.01.021

Newell, B. R., & Rakow, T. (2007). The role of experience in decisions from description. *Psychonomic Bulletin & Review, 14*(6), 1133-1133-1139.

Shanks, D. R., Tunney, R. J., & McCarthy, J. D. (2002). A re-examination of probability matching and rational choice. *Journal of Behavioral Decision Making, 15*(3), 233-233-250. doi:10.1002/bdm.413

Stanovich, K. E., & West, R. F. (2008). On the relative independence of thinking biases and cognitive ability. *Journal of Personality and Social Psychology, 94*(4), 672-672. doi:10.1037/0022-3514.94.4.672

Tversky, A., & Edwards, W. (1966). Information versus reward in binary choices. *Journal of Experimental Psychology, 71*(5), 680-683. doi:10.1037/h0023123

Tversky, A., & Kahneman, D. (1971). Belief in the law of small numbers. *Psychological Bulletin, 76*(2), 105-110. doi:10.1037/h0031322

Unturbe, J., & Corominas, J. (2007). Probability matching involves rule-generating ability: A neuropsychological mechanism dealing with probabilities. *Neuropsychology, 21*(5), 621-630. doi:10.1037/0894-4105.21.5.621

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

West, R. F., & Stanovich, K. E. (2003). Is probability matching smart? associations between probabilistic choices and cognitive ability. *Memory & Cognition, 31*(2), 243-251.

Wilson, W. A. (1960). Supplementary report: Two-choice behavior of monkeys. *Journal of Experimental Psychology, 59*(3), 207-207-208. doi:10.1037/h0041915

Wolford, G., Miller, M. B., & Gazzaniga, M. (2000). The left hemisphere's role in hypothesis formation. *Journal of Neuroscience, 20*(6), RC64.

Wolford, G., Newman, S. E., Miller, M. B., & Wig, G. S. (2004). Searching for patterns in random sequences. *Canadian Journal of Experimental Psychology/Revue Canadienne De Psychologie Expérimentale, 58*(4), 221-221. doi:10.1037/h0087446

Yellot, J. I. (1969). Probability learning with non-contingent success. *Journal of Mathematical Psychology, 6*, 541-575.