Memory & Cognition
2010, 38 (6), 667-676
doi:10.3758/MC.38.6.667

Probability matching and strategy availability

Derek J. Koehler and Greta James
University of Waterloo, Waterloo, Ontario, Canada

Findings from two experiments indicate that probability matching in sequential choice arises from an asymmetry in strategy availability: The matching strategy comes readily to mind, whereas a superior alternative strategy, maximizing, does not. First, compared with the minority who spontaneously engage in maximizing, the majority of participants endorse maximizing as superior to matching in a direct comparison when both strategies are described. Second, when the maximizing strategy is brought to their attention, more participants subsequently engage in maximizing. Third, matchers are more likely than maximizers to base decisions in other tasks on their initial intuitions, suggesting that they are more inclined to use a choice strategy that comes to mind quickly. These results indicate that a substantial subset of probability matchers are victims of "underthinking" rather than "overthinking": They fail to engage in sufficient deliberation to generate a superior alternative to the matching strategy that comes so readily to mind.

When faced with the task of choosing between two possible outcomes that differ in their probability of occurrence, many people match their choice probabilities to the corresponding outcome probabilities (matching) rather than always choosing the outcome with the higher probability (maximizing). For instance, in a task in which Outcome A occurs on 70% of the trials and Outcome B on the other 30%, probability matching involves predicting Outcome A on 70% of trials and B on the other 30%. By contrast, maximizing involves predicting the more likely Outcome A on every trial. Many people engage in matching even though maximizing is associated with a higher rate of predictive accuracy and, when participants are paid for each correct prediction, with a greater expected payoff. For a review of the literature on matching, which dates back to probability learning experiments in the 1950s, see Vulkan (2000).

Recently, researchers have adopted a dual-system approach to the study of probability matching (Kogler & Kuhberger, 2007; West & Stanovich, 2003). Dual-system accounts, particularly of the kind Evans (2008) calls "default interventionist," have become increasingly prevalent in the study of judgment, decision making, and reasoning. Probability matching, by such an account, reflects an intuitive response that arises from the operations of a fast, effortless, heuristic evaluation that could be, but often is not, overridden by a slower, more effortful deliberative assessment identifying maximizing as a superior alternative strategy. Consistent with this account is the finding that individuals who are higher in cognitive ability, and thereby presumably more proficient in deliberative reasoning, are more likely to maximize and less likely to probability match than are those of lower cognitive ability (Stanovich & West, 2008; West & Stanovich, 2003). Furthermore, manipulations that can be interpreted as encouraging deliberation, such as instructing participants to recommend a strategy to another person (Fantino & Esfandiari, 2002) or to think like a statistician (Kogler & Kuhberger, 2007), have been found to increase rates of maximizing behavior.

Koehler and James (2009) offered an elaboration of the dual-system account in which the initial intuition that gives rise to probability matching stems from relatively effortless operations that evaluate relevant outcome probabilities and use them to generate expectations regarding an upcoming sequence of events. Thus, as in the example above, in which Outcome A occurs on 70% of the trials and Outcome B on the other 30%, the intuitive system produces an expectation that, say, A will occur on 7 of the next 10 trials and B on the remaining 3. When faced with the task of determining how to make choices over those next 10 trials, then, the expectation generated by the intuitive system (expect 7 As and 3 Bs) serves as a natural candidate for making choices (predict 7 As and 3 Bs). This represents a form of "attribute substitution" (Kahneman & Frederick, 2002), in which the answer to a relatively difficult question (how many times should A and B be predicted?) is replaced by the answer to an easier one (how many times are A and B expected?).

Not everybody engages in probability matching, of course, suggesting that in some cases people are able to override the initial intuition that encourages matching, and instead use an alternative strategy that yields better returns (typically maximizing; in our studies, use of these two strategies constitutes the two modal responses). From the dual-system perspective, it is assumed that it is the deliberative system that overrides the initial, intuitive tendency to match in such circumstances. This override operation

D. J. Koehler, dkoehler@uwaterloo.ca

can fail for at least two reasons, with the consequence that the initial tendency to match is often "endorsed" in the final choices made by the decision maker. First, the deliberative system is assumed to require substantial cognitive resources to operate effectively. If those resources are not recruited for the choice task, the deliberative system may not be sufficiently engaged to override the intuitive matching response. Second, even if the deliberative system is fully engaged, if it does not have an alternative strategy (such as maximizing) in its repertoire (i.e., if that strategy is not already available in memory via recall or generation), it may have no recourse but to deploy the matching strategy.

In the present experiments, we investigate the role of strategy availability in determining the likelihood that a decision maker will engage in matching or maximizing. We hypothesize that probability matching arises from the unavailability of a superior alternative strategy—namely, maximizing. On this hypothesis, we would expect to see higher rates of endorsement and use of maximizing (and lower rates of matching) when the maximizing strategy is made highly available than when it is not.

Preliminary evidence for this claim was reported by Koehler and James (2009). Their participants completed a choice task in which they first learned some information about the contents of a bag of marbles—namely, that about 75% of the marbles in the bag were green and the other 25% were red—and then made guesses about which color would emerge on each of a series of draws (with replacement) from the bag. Participants received 50¢ for each correct guess. The majority of participants' choices followed a matching strategy (guessing "green" on 75% of draws and "red" on the other 25%), although a sizable minority of participants engaged in strict maximizing (guessing "green" for every draw). After completing the choice task, participants were presented with a number of strategies, including matching and maximizing, that might have been used during the game. 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. Once the maximizing strategy was described to them, the majority of participants (including roughly 40% of participants who had engaged in strict matching on the choice task) endorsed maximizing as superior to matching in terms of expected payoffs.

One interpretation of this finding, which we favor, is that many people match on the typical choice task because the matching strategy comes readily to mind (for the reasons described above), whereas maximizing does not. When both strategies are equated in terms of their availability in memory, as in the strategy comparison question posed to participants by Koehler and James (2009), maximizing is endorsed at a higher rate than it is used on the choice task. On this account, the difference in rates of matching versus maximizing between the choice task and the strategy comparison question is attributable to the (sizable) subset of participants who apparently match on the choice task because the maximizing strategy (which they later endorse as superior) is not available in memory at the time they make their choices. We refer to this as the "underthinking" account because it attributes matching behavior to a failure to retrieve or generate a strategy that, once available, would have been recognized as superior.

An alternative interpretation, which might be called the "overthinking" account, is that participants do in fact consider the maximizing strategy in the choice task but reject its use in favor of matching. On this account, there is no problem with unavailability of the maximizing strategy. Instead, the decision maker chooses to use a different strategy. This could be because maximizing is viewed as too simple to possibly be the correct strategy, or because it is too boring to implement, or because it is viewed as accepting the certainty of some of the predictions being wrong.[1] On this account, participants may be aware, when completing the choice task, that maximizing is superior in terms of expected returns but may choose not to use it because they are focused on other, or additional, considerations that lead them to engage in matching instead.

We conducted two experiments to disentangle these alternative interpretations. In the first, after completing a choice task in which they could use a maximizing or a matching strategy, participants were asked not only which strategy is superior in terms of expected payoffs, but also which they would use if they were to complete the choice task again. On the underthinking account, we would expect participants who match on the choice task simply because maximizing does not come to mind as an alternative strategy, not only to endorse maximizing as superior in terms of expected returns on the strategy comparison question, but also to say that they would maximize rather than match if they were to complete the choice task again. By contrast, on the overthinking account, participants may identify the maximizing strategy as having higher expected returns on the strategy comparison question, despite having matched on the choice task, because they are trying to accomplish something other than maximize expected returns. On this account, they should indicate that they would match again if they had another chance to complete the choice task. In the second experiment, we manipulated availability of the maximizing strategy prior to completion of the choice task. On the underthinking account, making the maximizing strategy more available should lead to its greater use on the choice task. On the overthinking account, by contrast, this manipulation ought to have no effect, because it is presumed that the maximizing strategy is already available even in the absence of this manipulation but that participants choose to use another strategy instead.

We also evaluate, in both experiments, how those who match and those who maximize on the choice task perform on other tasks that can be interpreted as measuring the extent to which an individual is inclined to base judgments and choices on intuition rather than deliberation. On the underthinking account, we would expect matchers—compared with maximizers—to give more intuitive and fewer deliberative responses on these tasks, indicating a deliberative system that is less effective in overriding initial intuitions that are incorrect, either because the deliberative system does not subject the intuitive response to much scrutiny or because it does not have a better solution in its repertoire. The role of such individual differences in cognitive ability and thinking dispositions as they pertain to judgment, decision making, and reasoning has been explored in extensive research programs by Stanovich and West (see, e.g., 2000) and by Cacioppo and Petty (see, e.g., 1982), among others.

EXPERIMENT 1

To minimize the possibility that participants engage in matching or some other suboptimal strategy simply because they are confused, we developed a simple, tangible version of the choice task. Participants were presented with 10 pairs of cups, each pair consisting of one red and one green cup. They were told that a $1 Canadian coin (referred to in Canada as a "loonie") had been hidden under either the green or the red cup in each pair. Participants were shown a 10-sided die, of which 7 sides had been painted green and the other 3 sides had been painted red.[2] For each pair of cups, participants were told that the experimenter had determined where to hide the coin by rolling the die, and then placing the coin under the green cup if a green side was rolled, or placing it under the red cup if a red side was rolled. Participants were asked, for each pair of cups, to drop a ring over one cup to indicate their guess as to the location of the loonie. At the end of the experiment, they were told, the cups would be lifted and they could keep all the loonies found under the cups on which they had dropped rings. Before the cups were lifted, participants completed several strategy comparison questions in which the matching and maximizing strategies were presented for direct evaluation.

This variant of the choice task has several notable features. First, participants do not need to learn the relevant outcome probabilities, which instead are precisely known from features of the 10-sided die (as in earlier work by Gal & Baron, 1996). Second, for a participant to engage in strict maximization requires only that the dominant outcome be chosen on a series of 10 trials, as opposed to other studies in which the required number is much higher and, presumably, more tedious to execute (e.g., Shanks, Tunney, & McCarthy, 2002, required that a participant choose the dominant option for 50 sequential trials before being considered to have maximized). Third, because we use a physical version of the task rather than a computer-based one, participants need not be concerned that they have been deceived regarding the outcomes (i.e., actual location of the coins). Fourth, the stakes per trial are nontrivial at Can$1 each, as opposed to a more extended trial sequence in which the stakes per trial would typically be substantially lower. Fifth, participants do not receive any kind of trial-by-trial outcome feedback at the time they make their choices, which minimizes concerns that (despite having been informed that the hiding places were randomly determined) they might seek patterns in the outcome sequence or modify their choices on the basis of previous outcomes (e.g., in accord with the gambler's fallacy). Sixth, outcome feedback is withheld until participants have completed the strategy comparison questions, so that they have no additional information (e.g., the observation that using the matching strategy did not pay well) at the time they complete those questions.

Our choice task differs in some potentially significant ways from the typical probability learning task in which matching behavior was first studied. In those tasks, predictions are made one at a time and are followed immediately with information about the outcome, and typically, the participant must rely on memory for a record of previous predictions and their associated outcomes. In our task, by contrast, predictions are made for a set of simultaneously visible choice options, and outcome feedback is withheld until all predictions have been made. The standard probability learning paradigm may give rise to certain kinds of reasoning (e.g., a search for patterns, or belief in the gambler's fallacy) that our task effectively blocks. The role of strategy availability in probability matching behavior in the standard probability learning paradigm is something we plan to explore in future research.

Method

Participants. Eighty-four students (49 male, 29 female, 6 not recorded) were recruited from a campus student center and were paid up to Can$10 for their participation, depending on the accuracy of their performance.

Apparatus. As described above, the task involved 10 pairs of cups, with one red and one green cup in each pair. In one variant of the task, a Can$1 coin was hidden under either the red or green cup in each pair. In another variant, paper checkmarks were hidden instead, which were redeemed for Can$1 each at the end of the experiment. These two variants produced equivalent results, so this factor is not considered further in the analyses below. For each participant, the hiding place of the coins or checkmarks was determined for each pair of cups by the roll of a 10-sided die with 7 green faces and 3 red faces (or vice versa; the dominant color was counterbalanced across participants). Participants placed a black ring around either the red or green cup in each pair to indicate their prediction that the item was hidden under that cup. The pairs of cups were placed on a table in a tent that hid the game and outcomes from passersby.

Procedure. The experimenter explained the game to participants as they sat in front of the table on which the pairs of cups had been placed. They were told that the coins or checkmarks had been hidden, as described above, prior to their entering the tent. Participants were told that they would receive Can$1 for each correct guess. After participants had made all 10 choices by placing the 10 rings, but before the cups were lifted, participants completed a brief questionnaire. This questionnaire described the strategies used by two hypothetical characters, Mike and John. Mike had guessed green for all 10 pairs of cups; John had guessed green for 7 pairs and red for the other 3 pairs. Thus, Mike and John used the maximizing and matching strategies, respectively, although no strategy labels were given nor was one strategy described as better than the other. Participants were asked to indicate whose strategy, Mike's or John's, (1) their own strategy more closely resembled, (2) would be expected to earn more money, and (3) 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.

After playing the game, participants completed the Cognitive Reflection Test (CRT; Frederick, 2005), which consists of three problems. Each problem typically elicits a potential answer that comes quickly to mind but is in fact incorrect. The first problem, for example, is "A bat and a ball cost $1.10. The bat costs $1.00 more than the ball. How much does the ball cost?" A possible answer, 10¢,

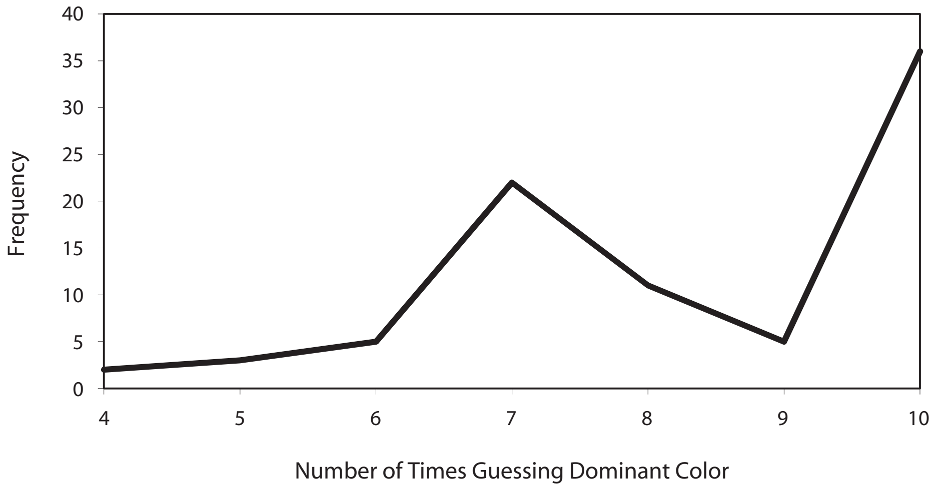

Figure 1. Distribution of number of dominant color choices in Experiment 1, where the expected number of such choices is 7, given a matching strategy, and 10, given a maximizing strategy.

comes quickly to most people's minds but can be readily determined to be incorrect. The correct answer, 5¢, seems to require a bit more thought. Thus, the CRT can be interpreted as a measure of "cognitive reflection—the ability or disposition to resist reporting the response that first comes to mind" (Frederick, 2005). Participants also completed 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

Figure 1 shows the distribution of the number of dominant color guesses made by participants on the choice task. Of the 84 participants, 36 (43%) engaged in strict maximization (dominant color guesses = 10). Another 22 (26%) engaged in strict matching (dominant color guesses = 7). If we categorize participants whose number of dominant color guesses falls within 1 of what is expected under maximizing as maximizers ($n = 41$) and those whose number of dominant color guesses falls within 1 of what is expected under matching as matchers ($n = 38$), and ignore the small number of participants ($n = 5$) who chose the dominant color on 5 or fewer trials, then approximately half the participants are classified as maximizers and the other half as matchers.

We now compare choice task performance with responses to the three strategy comparison questions (where, due to a few missing answers, $N = 83$, 81, 83 for self-reported strategy use, strategy with highest expected payoff, and intended strategy if the game were to be played again, respectively). Answers to the first question, in which participants identified which strategy they had used on the choice task, coincided closely with our classification based on number of dominant color guesses, as is shown in Figure 2. In contrast to both the observed and self-reported strategies used on the choice task, once both strategies were presented to participants, the large majority endorsed maximizing as superior to matching in terms of expected returns and also indicated that they would use the maximizing strategy were they to play the game again (Figure 2). Compared with the proportion of participants who engaged in strict maximizing on the choice task, a significantly larger proportion endorsed maximizing as superior in terms of expected returns [$\chi^2(1, N = 81) = 17.1, p < .001$] and as the strategy they would use if they were to play the game again [$\chi^2(1, N = 83) = 16.9, p < .001$]. Even among those who engaged in strict matching on the choice task, in the direct comparison between the two strategies, 9 of 21 (43%; 1 participant did not answer) endorsed maximizing as superior to matching in terms of expected returns, and the same proportion of participants said they would maximize if they were to play again.

Next, we turn to scores on the CRT, which indicate the number of correctly answered problems and therefore range from 0 to 3. The correlation between number of dominant color guesses on the choice task and scores on the CRT is .40 ($p < .001$), indicating that, compared with low scorers, high scorers on the CRT are more likely to maximize and less likely to match on the choice task. Figure 3 plots mean number of dominant color choices by CRT score and indicates that those scoring 0 on the CRT largely match, whereas those scoring 3 approach strict maximizing. The relation between CRT score and maximizing remains statistically significant even after controlling for mathematical ability, as measured either by self-reported proficiency or number of math courses taken.

We also examined how responses to the strategy comparison questions were related to CRT scores. When participants were asked to identify the strategy they had used in the choice task, the correlation of their responses with CRT scores was nearly identical ($r = .41, p < .001$) to that between actual choices and CRT scores reported above. But correlations with CRT scores were lower for the questions asking which strategy has higher expected payoffs and which strategy participants would use if they were to play again ($r = .22$ and .19, $p = .05$ and .08, respectively). This difference appears attributable to the higher rate of endorsement of the maximizing strategy when it is made available, in comparison with its spontaneous use in the choice task, a difference that is larger for low than for high

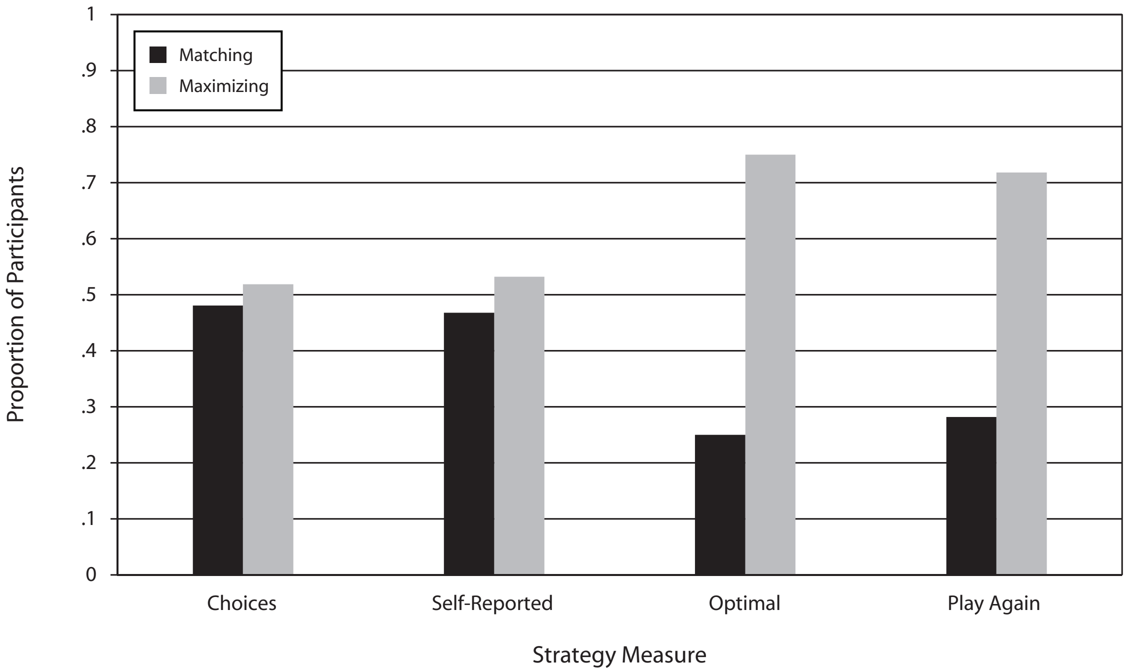

Figure 2. Proportions of participants in Experiment 1 selecting a matching and maximizing strategy (1) in their actual choices on the choice task, (2) in their self-reported strategy use, (3) in their judgments of which strategy would be expected to produce higher payoffs, and (4) in their reports as to which strategy they would use in the choice task were they to play again. Excluded from the figure are data from 5 participants who chose the dominant color on the choice task on only 50% or fewer of their choices.

CRT scorers. As an illustration, consider the difference between the two measures for the lowest CRT scorers (score of 0, $n = 18$) and the highest scorers (score = 3, $n = 23$): In their choice behavior, only 2 of the 18 low scorers (11%) engaged in strict maximizing, compared with 17 of 23 high scorers (74%). But 11 of the 18 low scorers (61%) endorsed maximizing as superior in expected returns to matching, as did 19 of the 21 high scorers (91%; 2 participants did not answer this question). This result suggests that the CRT better distinguishes spontaneous strategy use in the choice task than it does strategy endorsement when both strategies are made available following completion of the choice task.

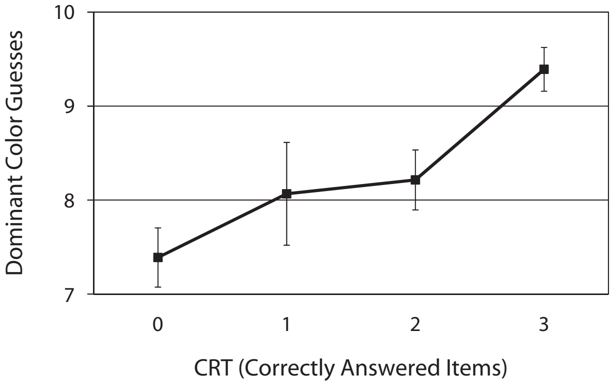

Figure 3. Mean number of dominant color choices by Cognitive Reflection Test (CRT) score in Experiment 1. Numbers of participants scoring 0, 1, 2, and 3 on the CRT were 18, 15, 28, and 23, respectively. Error bars indicate standard errors of the means.

Discussion

Two results from Experiment 1 support the "underthinking" account over the "overthinking" account of probability matching behavior. First, compared with the rate of maximizing on the choice task, many more participants endorsed the maximizing strategy as superior to matching and said that they would use it if they were to play again when both strategies were made available in the strategy comparison questions. Second, compared with low scorers, high scorers on the CRT, for which good performance requires the deliberative system to override an incorrect initial intuition, were more likely to maximize on the choice task.

EXPERIMENT 2

Participants in the previous experiment indicated in the direct strategy comparison that they would engage in maximizing rather than matching if they were to play the game again. Experiment 2 was designed to corroborate such claims by manipulating the availability of the maximizing strategy prior to the choice task. On the underthinking account, making the maximizing strategy more highly available should lead to its greater use in the subsequent choice task. On the overthinking account, by contrast, this manipulation should have no effect, because it is assumed that participants readily generate maximizing as a potential strategy but choose not to use it.

Method

With a few small exceptions noted below, the task and procedure were identical to those of the previous experiment. The only critical difference was that one group of participants, randomly assigned to the "hint" condition, were presented with both potential strategies prior to the choice task and were asked which one would be expected to earn them more money, as follows:

> 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 7 sets of cups and red for 3 sets of cups.
>
> Which strategy, 1 or 2, do you think will win you the most money?

Participants in the no-hint condition did not answer this question before completing the choice task. Following completion of the choice task, all participants were presented with the two strategies as described above and were asked which strategy, 1 or 2, (1) their own strategy more closely resembled, (2) would be expected to earn more money, and (3) they would use if they were to play the game again.

In addition to the CRT, which was administered at the end of the session just before the demographic and math skills questions, participants also completed another task thought to gauge proneness to reliance on intuition versus deliberation, which we will refer to as the urns task. It is adapted from Epstein's jelly beans task (see, e.g., Denes-Raj & Epstein, 1994). Participants are asked to consider two urns. One contains 1 gold ball and 9 white balls; the other contains 9 gold balls and 91 white balls. Participants are asked to imagine that they are to make a single draw, at random, from one of the urns and are told that if they draw a gold ball, they will win a free vacation. Although the urn with only a single gold ball offers the higher probability of winning, many people experience and even express a preference for the urn that offers the larger absolute number of gold balls. Participants were asked to give ratings on which urn (1) offered the higher probability of drawing a gold ball, (2) they felt would be easier to win with when they drew from it, (3) would be more exciting to draw from, (4) they would choose to draw from, and (5) they would pay more to draw from. Ratings were made on a 5-point scale (1 = *definitely Urn A*; 3 = *Urn A and Urn B are equal*; 5 = *definitely Urn B*—where Urn A was the one with 1 gold and 9 white balls). Epstein and his colleagues depicted the jelly beans task as putting the intuitive and deliberative systems into conflict. Pacini and Epstein (1999) reported that the higher probability option is chosen more often by high scorers than by low scorers on the rational thinking component of the Rational–Experiential Inventory, a measure of individual differences in thinking styles.

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 extra course credit and played a low-stakes version of the game in which each correct guess paid Can$0.25. In the second variant, conducted in the same public location as Experiment 1, university students ($N = 98$; 40 females) played a high-stakes version of the game in which they could win up to Can$10 as compensation for their participation; in this version of the game, each correct guess paid Can$1.

Results

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

In the same ANOVA, condition (hint vs. no hint) had a significant main effect on number of dominant color guesses [$F(1,154) = 7.17$, $p = .008$, $\eta_p^2 = .044$]. In the no-hint 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 4 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). By contrast, among participants in the hint condition, the rate of strict

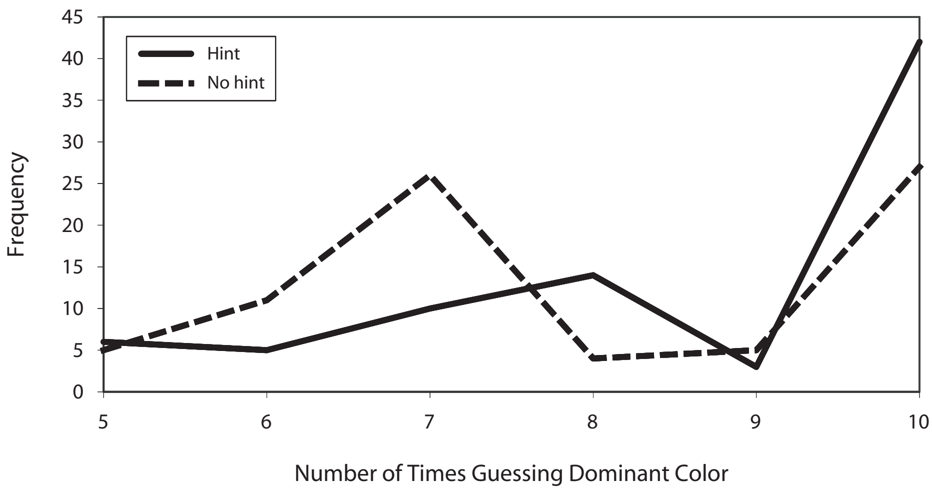

Figure 4. Distribution of number of dominant color choices in Experiment 2, where the expected number of such choices is 7, given a matching strategy, and 10, given a maximizing strategy, for the hint (solid line) and no-hint (dashed line) conditions.

maximization was significantly higher, at 53% (42 out of 80 participants) [$\chi^2(1, N = 158) = 5.14, p = .023$]. The hint manipulation had a marginally significant impact on self-reported strategy use as elicited following the choice task [$\chi^2(1, N = 158) = 3.05, p = .081$], 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.

The hint manipulation did not have a significant effect on the other strategy questions asking which had 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 (1 person did not complete this item) and 53 of 78 participants in the no-hint condition said that 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 5 shows, for each condition, the proportion of participants classified as matchers or maximizers on the choice task (using the same criteria as in Experiment 1) 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.[3] 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 posttask 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 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) [$\chi^2(1, N = 123) = 7.99, p = .005$].

CRT scores once again correlated positively with number of dominant color guesses ($r = .36, p < .001$). On the urns task, judgments favoring the urn with the larger absolute number but a smaller proportion of winning balls correlated negatively with number of dominant color guesses. Ratings for four of the five questions about the urns (all but the one asking which was more exciting to draw from) exhibited statistically significant correlations in this direction; the correlation between the sum of the five ratings and number of dominant color guesses was $-.33$ ($p < .001$). Both CRT and the sum of ratings on the urns task remain statistically significant when entered simultaneously as predictors with number of math courses taken.

An interesting, if unexpected, result was that the hint manipulation had a significant effect on responses to the urns task, with those who received the hint more strongly favoring the urn that offered the higher probability of winning despite having a smaller absolute number of winning balls. The mean rating across the five items, for instance, on a 5-point scale in which lower scores indicated a stronger preference for the high-probability urn, was 2.1 ($SD = 0.9$) in the hint condition and 2.7 ($SD = 1.0$) in the no-hint condition [$t(155) = 3.74, p < .001$]. Scores on each individual item on the urns task were significantly lower in the hint than in the no-hint condition, with the exception of the item concerning which urn was more exciting to draw from. The hint manipulation did not significantly influence CRT scores.

Finally, the correlations between CRT and urns task performance with number of dominant color guesses on the choice task were computed separately for the hint and the no-hint conditions. Although the 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, correlations with number of dominant color choices were .43 for CRT and $-.38$ for the urns task (total of five items); in the no-hint condition, the corresponding correlations were .28 for CRT and $-.23$ for the urns task. Compared with participants prone toward reliance on intuition, those prone toward reliance on deliberation seemed to benefit more from the hint. Figure 6 displays this pattern of results for the CRT scores. Notably, those scoring 0 on the CRT (i.e., those presumed to be most reliant on intuition) apparently did not benefit from the hint at all.

We tried to gain further insight into this result by examining how responses to the strategy comparison questions vary with CRT scores. In answer to the pre-choice-task question as to which strategy had higher expected payoffs (i.e., the "hint" that was intended to make the maximizing strategy more readily available), 24 of the 28 participants (86%) with CRT scores of 3 selected the maximizing strategy; by comparison, only 8 of the 18 participants (44%) with CRT scores of 0 did so. Among these same low CRT scorers who received the hint, furthermore, experience on the choice task did not appear to provide any additional insights as to the benefits of maximizing, as shown by their posttask responses to the strategy comparison questions: Only 7 of 18 (39%) endorsed maximizing as having higher expected payoffs (all of whom had also done so before the choice task), and only 6 (33%) said that they would use the maximizing strategy if they were to play again (including 4 who had endorsed maximizing and 2 who had endorsed matching before the choice task). Interestingly, those with CRT scores of 0 who did not receive the hint were more likely, after completing the choice task, to endorse the maximizing strategy as superior in terms of expected payoffs (15 of 21 participants, or 71%) and to say that they would maximize if they were to play again (13 participants, or 62%). The sample sizes

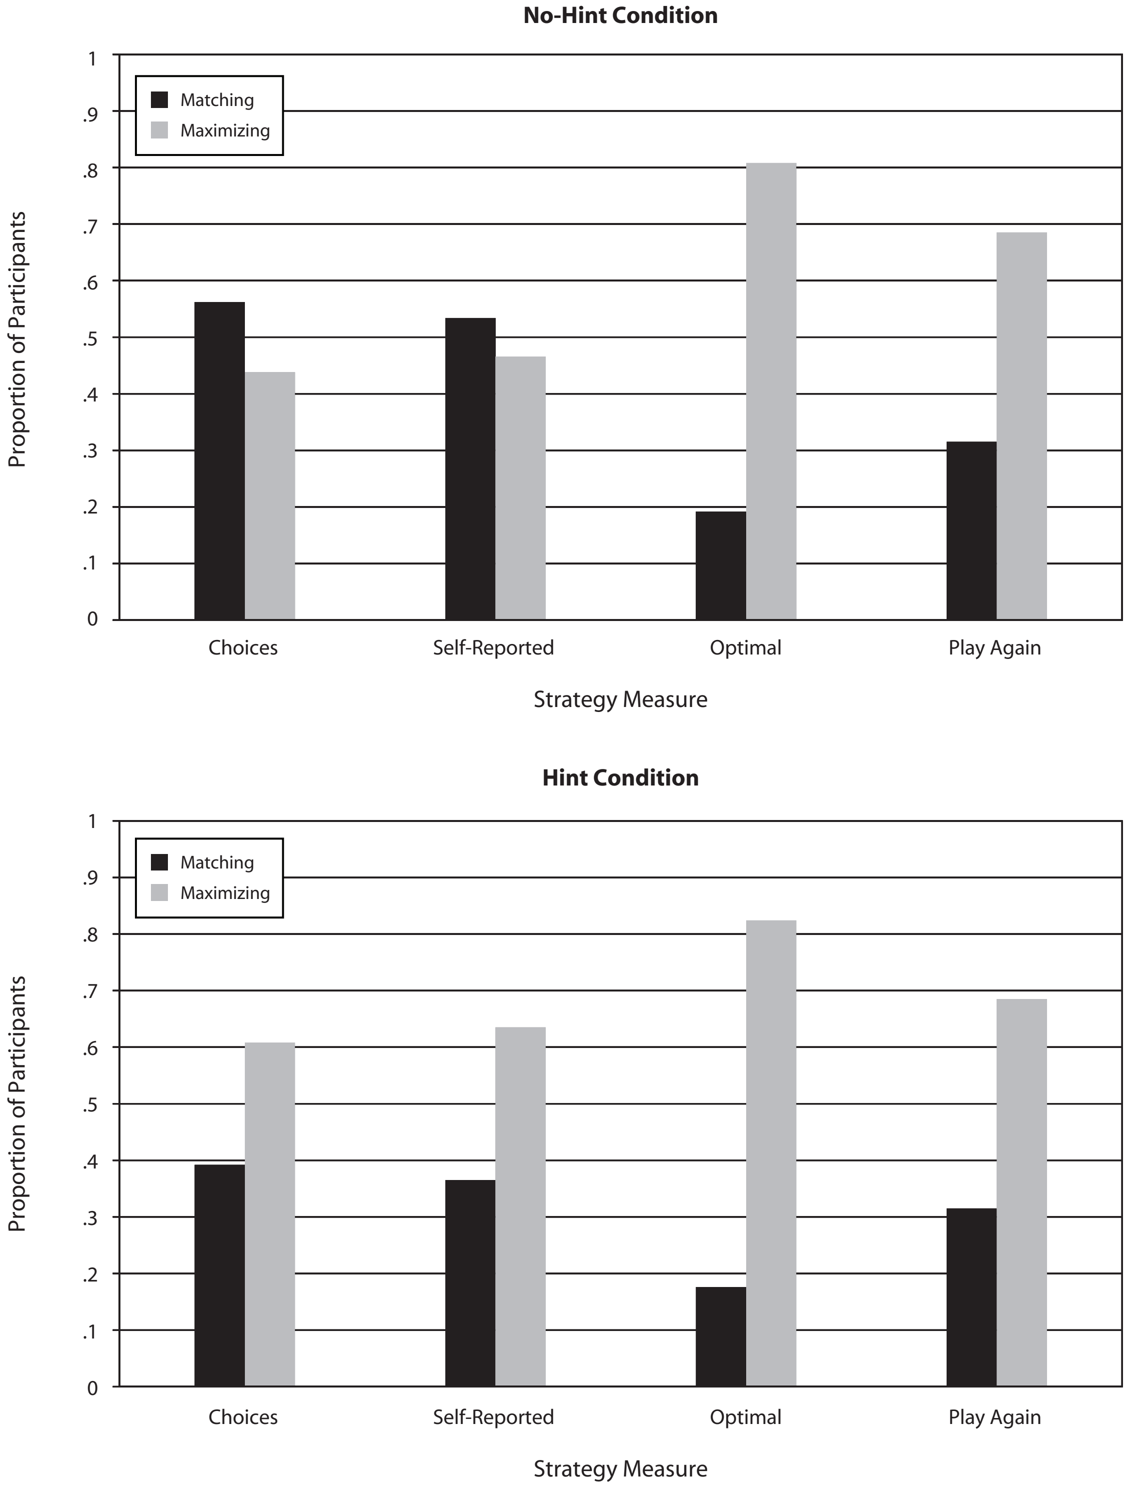

Figure 5. Strategy use and endorsement rates in Experiment 2, by condition. Excluded from the figure are data from 11 participants who chose the dominant color on the choice task on only 50% or fewer of their choices.

in these analyses are quite small, and the discrepancy in endorsement rates on posttask strategy comparison questions between those in the hint and no-hint conditions was limited to the very lowest scoring group on the CRT, so this result should be interpreted with caution. That said, it is interesting that both here and in Experiment 1, the majority of the lowest CRT scorers endorsed maximizing as superior when presented with that strategy explicitly for the first time after completion of the choice task, but not when presented with that strategy before completing the

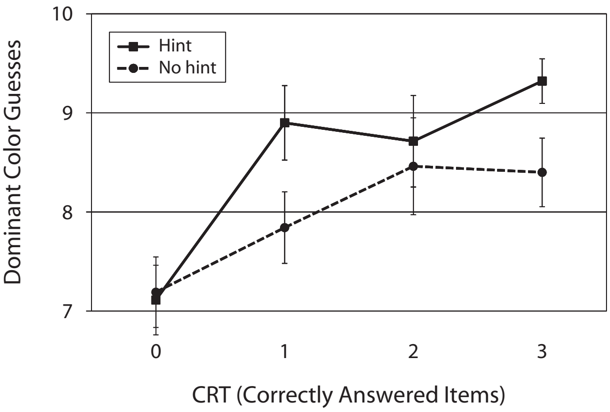

Figure 6. Mean number of dominant color choices by CRT score in Experiment 2, for the hint (solid line) and no-hint (dashed line) conditions. Numbers of participants scoring 0, 1, 2, and 3 on the Cognitive Recognition Test (CRT) were 18, 20, 14, and 28 in the hint condition and 21, 19, 13, and 25 in the no-hint condition, respectively. Error bars indicate standard errors of the means.

choice task. It is possible that this group more readily recognized the superiority of the maximizing strategy after having direct experience with the choice task.

GENERAL DISCUSSION

The results of the two experiments reported here support the claim that probability matching in choice arises, at least in part, from an asymmetry in strategy availability: The matching strategy comes readily to mind, whereas the maximizing strategy does not, and many people apparently engage in matching because they fail to generate a better strategy. Two findings from the present research directly support this account. First, many more participants endorse maximizing as superior to matching in a direct comparison when both strategies are made available than actually engage in maximizing in the choice task when they must generate the maximizing strategy on their own. Second, when the maximizing strategy is made more readily available prior to completion of the choice task, more participants subsequently engage in maximizing behavior.

A third major finding of interest is that, compared with those who match on the choice task, those who maximize score higher on the CRT (Experiments 1 and 2) and the urns task (Experiment 2). Both tasks place the intuitive and deliberative systems in potential conflict, and we take performance on these tasks to reflect chronic individual differences in the tendency for one or the other system to predominate. It should be acknowledged that exactly what these tasks measure is not completely known. In particular, there is some question about whether the CRT (and possibly the urns task as well) should be viewed as a measure of thinking dispositions (i.e., a preference for deliberation over intuition) or of cognitive ability (e.g., mathematical reasoning or even general intelligence). We did find that CRT scores continued to be predictive of choice task performance even when controlling for mathematical ability in our analyses, but our measures of such ability are certainly far from perfect. In some sense, either interpretation is compatible with the dual-system account we develop here, in that the initial intuition to match is expected to be overridden in favor of maximizing only if the individual is inclined (has a thinking disposition) to scrutinize the initial intuition and is able (has the cognitive ability) to generate a superior alternative strategy. More research is needed to address the interplay of cognitive ability and thinking dispositions as determinants of the rationality of human decisions (Stanovich, 2009).

In summary, then, we take our results to collectively suggest that a substantial subset of probability matchers on the choice task are victims of "underthinking" rather than "overthinking": They are unwilling or unable to generate a superior alternative to the matching strategy that comes so readily to mind. This result does not, of course, rule out the possibility that matching behavior in some cases arises from overthinking. Indeed, in Experiment 2, only about 70% of participants in the hint condition who correctly identified maximizing as superior to matching in terms of expected returns went on to engage in strict maximizing behavior on the subsequent choice task. In other words, approximately 30% of participants who had just identified maximizing as the superior strategy in terms of expected returns chose not to maximize on the choice task. It is possible that these participants were interested in maximizing something other than expected returns. Perhaps they were striving for perfect predictive accuracy and believed, erroneously, that maximizing was not the best way to achieve it. Or perhaps they felt that maximizing was too simple a strategy to be the optimal solution to the decision task they faced. It has also been suggested, although it does not apply to the context of our experiments, that matching may arise from a (futile) search for patterns in the random outcome sequence (see, e.g., Gaissmaier & Schooler, 2008). All of these possibilities result from what might be considered forms of overthinking rather than underthinking.

The results of Experiment 2 might be used as a crude means of identifying and estimating the relative predominance of different types of thinkers in the choice task. Roughly 20% of the participants (call them the "intuitionists") failed to identify maximizing as superior in the direct comparison question, instead endorsing matching as likely to produce higher returns. Approximately 40% of participants (call them the "rationalists") engaged in strict maximizing on the choice task even without the benefit of the hint manipulation. This leaves another 40% of participants (call them the "swayable" and the "conflicted") who would not maximize in the absence of a hint because they did not spontaneously generate the maximizing strategy, but who did endorse maximizing as superior to matching in terms of expected returns when the two strategies were presented to them for comparison. The overall level of maximizing behavior (about 60%) following the hint making that strategy readily available implies that these last two groups are of approximately equal size (i.e., each represents about 20% of our participant population), where the swayable group adapts its choice behavior to fit the strategy they have just identified as superior and the conflicted group persists in matching despite having identified maximizing as superior in the direct strategy comparison. Our analysis is helpful in decomposing participants into these groups, and our findings support an important role of strategy availability among one of the resulting subgroups—namely, the swayable. And the results of our individual difference analyses support the idea that there are fairly broad and chronic differences in ability and inclination toward deliberation between the intuitionists and the rationalists that could justify use of those labels. The beliefs and motivations of the conflicted group, however, at least in our work, remain largely unexplained.

AUTHOR NOTE

This research was funded by a grant from the Natural Sciences and Engineering Research Council of Canada to the first author. We are grateful to Lei Mangel and Julie Scott for their assistance with the experiments, and to seminar participants at the University of Chicago, the University of Pennsylvania, and the University of Waterloo for their helpful comments on this line of research. Address all correspondence to D. J. Koehler, Department of Psychology, University of Waterloo, Waterloo, ON, N2L 3G1 Canada (e-mail: dkoehler@uwaterloo.ca).

REFERENCES

Cacioppo, J. T., & Petty, R. E. (1982). The need for cognition. *Journal of Personality & Social Psychology*, **42**, 116-131.

Denes-Raj, V., & Epstein, S. (1994). Conflict between intuitive and rational processing: When people behave against their better judgment. *Journal of Personality & Social Psychology*, **66**, 819-829.

Evans, J. St. B. T. (2008). Dual-processing accounts of reasoning, judgment, and social cognition. *Annual Review of Psychology*, **59**, 255-278.

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

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

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

Gal, I., & Baron, J. (1996). Understanding repeated simple choices. *Thinking & Reasoning*, **2**, 81-98.

Kahneman, D., & Frederick, S. (2002). Representativeness revisited: Attribute substitution in intuitive judgment. In T. Gilovich, D. Griffin, & D. Kahneman (Eds.), *Heuristics and biases: The psychology of intuitive judgment* (pp. 49-81). New York: Cambridge University Press.

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

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

Pacini, R., & Epstein, S. (1999). The relation of rational and experiential information processing styles to personality, basic beliefs, and the ratio-bias phenomenon. *Journal of Personality & Social Psychology*, **76**, 972-987.

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**, 233-250.

Stanovich, K. E. (2009). Distinguishing the reflective, algorithmic, and autonomous minds: Is it time for a tri-process theory? In J. Evans & K. Frankish (Eds.), *In two minds: Dual processes and beyond* (pp. 55-88). Oxford: Oxford University Press.

Stanovich, K. E., & West, R. F. (2000). Individual differences in reasoning: Implications for the rationality debate? *Behavioral & Brain Sciences*, **23**, 645-665.

Stanovich, K. E., & West, R. F. (2008). On the relative independence of thinking biases and cognitive ability. *Journal of Personality & Social Psychology*, **94**, 672-695.

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

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

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

NOTES

1. In fact, in our tasks, maximizing offers not only the highest average return, but also the highest probability of guessing all the outcomes correctly; but if respondents follow the "law of small numbers" (Tversky & Kahneman, 1971) and anticipate that outcome proportions over the sequence will exactly equal their expected values, they might feel that maximizing entails accepting incorrect guesses that are avoidable only by matching.

2. For ease of exposition, we refer to green as the dominant outcome, but in fact whether the dominant color was red or green was counterbalanced across participants.

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

(Manuscript received November 5, 2009;
revision accepted for publication February 18, 2010.)