CHAPTER THREE

Probability Matching, Fast and Slow

Derek J. Koehler[1], Greta James
Department of Psychology, University of Waterloo, Waterloo, Ontario, Canada
[1]Corresponding author: e-mail address: dkoehler@uwaterloo.ca

Contents

Abstract

A prominent point of contention among researchers regarding the interpretation of probability-matching behavior is whether it represents a cognitively sophisticated, adaptive response to the inherent uncertainty of the tasks or settings in which it is observed, or whether instead it represents a fundamental shortcoming in the heuristics that support and guide human decision making. Put crudely, researchers disagree on whether probability matching is "smart" or "dumb." Here, we consider evidence for both "smart" and "dumb" variants of probability-matching behavior, as well as its alternative, maximizing. We rely on the influential and often-cited distinction between two "systems" of thinking to organize the research and competing interpretations of probability-matching behavior as "smart" or "dumb."

1. INTRODUCTION

Consider a simple computer game in which, on each trial, either a green or a red light appears. Your task is to predict which color will appear, and you will be paid a small amount of money for each correct prediction. What should you do, assuming your goal is to earn as much money as possible? Much of the challenge in this task arises from uncertainty regarding the

Psychology of Learning and Motivation, Volume 61
ISSN 0079-7421
http://dx.doi.org/10.1016/B978-0-12-800283-4.00003-4

process that determines whether the green or the red light appears on each trial (e.g., Green, Benson, Kersten, & Schrater, 2010). Does one light appear more frequently than the other? Is there a predictable pattern in the sequence of red and green outcomes? Does the probability of the green light illuminating change over the course of the game? Is it affected by your own actions, that is, the guesses that you have made on previous trials?

In the first experiments that investigated this type of task, one particular regularity in people's responses became the focus of researchers' attention: people tended to make their predictions in a manner that matched the relevant outcome probabilities (Goodnow, 1955; Grant, Hake, & Hornseth, 1951). For instance, if the green light was illuminated on 70% of trials and the red light on the remaining 30%, people tended to predict green on 70% of the trials and red on the remaining 30%. This phenomenon is referred to as *probability matching*, which can be defined more generally as the tendency to match choice proportions to outcome proportions in a binary prediction task.

In the experiments investigating probability matching that are of interest here, the outcomes being predicted (e.g., green vs. red light) are determined by a random process that is serially independent and stationary, which means that the probability of, say, the green light illuminating is the same on every trial, regardless of what occurred on the previous trial or how many trials have elapsed. Under such circumstances, it is easy to show that, if one's goal is to maximize the number of correct predictions, probability matching is inferior to an alternative strategy in which the higher probability outcome is predicted on every trial. This superior strategy is referred to as *maximizing*. For example, when the probability of a green outcome is 70%, maximizing (predicting green on every trial) yields an average predictive accuracy of 70%, while matching (predicting green on 70% of trials and red on the other 30%) yields an average predictive accuracy of $(0.70 \times 0.70) + (0.30 \times 0.30) = 58\%$.

Probability matching has attracted interest because it represents a violation of a cornerstone principle of rational choice theory, referred to as stochastic dominance. According to this principle, a gamble offering a probability P of some desired outcome should always be preferred to an otherwise equivalent gamble offering a lower probability P^* of obtaining that same outcome. Under conditions where payment is received for each correct prediction, a probability matcher violates the principle of stochastic dominance every time he or she predicts the lower probability outcome (e.g., red in the example above in which red occurs on only 30% of trials). In other words, he or she is choosing a gamble with a lower probability P^* of receiving a payment over one that offers a higher probability P of receiving the same payment. Probability matching, in short, appears anomalous from the perspective of rational choice models and for that reason demands explanation.

What causes probability matching? Who does it, and under what circumstances is it more or less likely to occur? In this chapter, we review some possible answers to these questions that have been offered in recent research on the topic. Our intent is not to systematically and exhaustively review every article that has been published on probability matching. Our review is highly selective and focuses almost exclusively on what might be referred as a "second wave" of research on probability matching that has taken place over the past decade or so. The voluminous original work on the topic in the 1950s and 1960s is not reviewed here; instead, the reader is directed to a helpful review by Vulkan (2000). It is worth noting, however, that several features of that early work, rooted almost exclusively in the then-predominant probability-learning paradigm, have proved unnecessary to observe probability-matching behavior. For instance, as is elaborated later in this chapter, probability matching is observed even in tasks in which the relevant outcome probabilities are known to participants from the outset, rather than having to be learned via trial-by-trial outcome feedback (e.g., Gal & Baron, 1996). In other words, even when many of the questions a person might have about the binary prediction task, such as those in our opening paragraph, are circumvented, probability matching is still regularly observed.

The most prominent point of contention among researchers regarding the interpretation of probability-matching behavior is whether it represents a cognitively sophisticated, adaptive response to the inherent uncertainty of the tasks or settings in which it is observed, or whether instead it represents a fundamental shortcoming in the heuristics that support and guide human decision making. Put bluntly, researchers disagree on whether probability matching is "smart" or "dumb." Our use of these terms is not intended to be entirely pejorative. Rather, they can be used to characterize, for instance, a person's own response—after having engaged in probability matching—to the argument that maximizing is a superior strategy. The person might explain, based on their understanding of the task, why it might have been reasonable to engage in matching (i.e., that it was a "smart" response to the task); alternatively, the person might do a forehead slap and acknowledge that they made a mistake (i.e., that matching was a "dumb" response to the task). Researchers themselves have disagreed as to whether or not probability matching should be viewed as a mistake. One of our main goals in this chapter is to organize the theoretical depictions of probability matching that have been offered around this admittedly crude distinction between "smart matching" and "dumb matching" accounts.

We rely on the influential and often-cited distinction between two "systems" of thinking (e.g., Kahneman & Frederick, 2002; Sloman, 1996; Stanovich & West, 2000; for a review, see Evans, 2008) to organize the research and competing interpretations of probability-matching behavior. One category of cognitive and affective processes shares the characteristics of being fast, effortless, unintentional, and unavailable to conscious awareness; the other category is relatively slow, effortful, intentional, and available to conscious awareness. We refer to the former as constituting the "intuitive" system and the latter as the "deliberative" system. We also adopt Kahneman's (2011; Kahneman & Frederick, 2002) characterization of the relation between the two systems as one in which the output of the intuitive system is imperfectly monitored and sometimes corrected or overridden by the deliberative system. In particular, Kahneman's account identifies well-known judgmental heuristics with the operations of the intuitive system and attributes many biases of judgment to a substitution process in which a person faced with a particular question receives from the intuitive system the answer to a different question but fails to recognize the discrepancy, and instead "endorses" that answer. This process of "attribute substitution" is discussed further below as it pertains to probability matching.

From a dual-system perspective, then, probability matching is "dumb" when it emerges from an intuitive response to the prediction task that goes uncorrected by the deliberative system. Maximizing, by this account, is "smart" when it results from the deliberative system correcting or overriding the intuition that makes matching compelling. Conversely, there may be circumstances under which maximizing represents the intuitive system's initial response to the task, giving rise to "dumb" maximizing. By contrast, probability matching under certain circumstances may emerge as the product of effortful deliberation (e.g., in which maximizing is considered but rejected as a possible strategy), which would be a case of "smart" matching. The remainder of the chapter organizes the recent findings of studies on probability matching in terms of evidence supporting smart and dumb variants of probability-matching and -maximizing behavior.

2. DUMB MATCHING

We will use a task, developed by Koehler and James (2010), as a running example of the characterization of probability matching as a "dumb" or intuitive response. As shown in Fig. 3.1, participants were presented with 10 pairs of cups, placed upside down on a table. Each pair consisted of one green and one red cup. Participants were told that, before they had entered the room, a dollar coin had been hidden under one member of each pair of cups. Which cup in the pair, green or red, the coin had been placed under had been determined by the roll of a 10-sided die with seven green faces and three red faces.[1] Participants were instructed to guess, for each pair, under which cup the coin was hidden by dropping a black ring over the cup. Participants were informed that once all 10 guesses had been made, the cups would be turned over and the participant could keep all the coins whose location had been correctly predicted.

Probability matching in this task would entail making seven green and three red predictions. Maximizing would entail making 10 green predictions. Of course, participants did not limit themselves to these two strategies, but matching and maximizing did represent the two modal responses to the task, as is shown in Fig. 3.2. In this particular study, fewer participants engaged in matching than in maximizing; in other studies, we have found the opposite. For present purposes, the important observation is that matching and maximizing emerged as two commonly used strategies in the cups task.

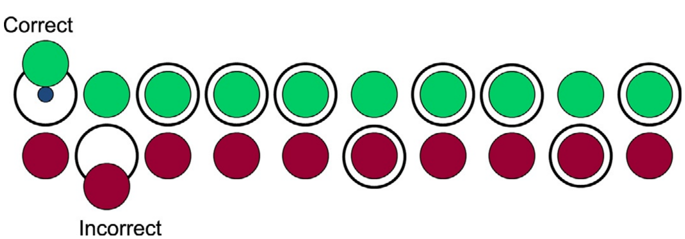

Figure 3.1 Schematic illustration of the cups task developed by Koehler and James (2010). In this example, the participant's predictions (indicated by black rings) follow a probability-matching strategy.

[1] For ease of discussion, we will refer to green as being the more probable color. In fact, this variable was counterbalanced such that, for half of the participants, red was the more likely outcome.

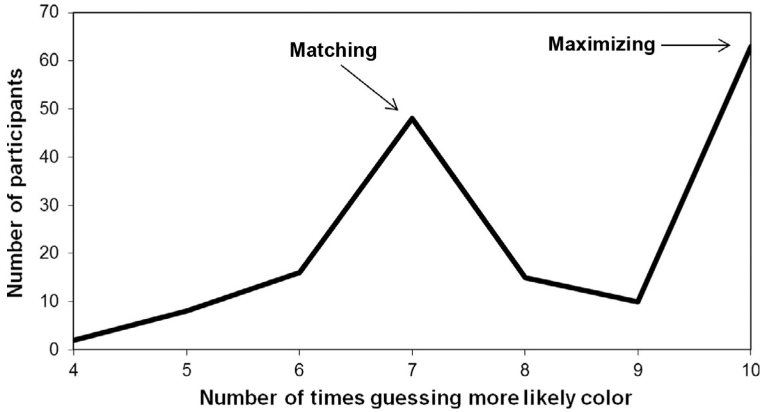

Figure 3.2 Number of times the more likely outcome was predicted, out of a possible 10, in Koehler and James (2010; collapsed across Study 1 and the no-hint condition of Study 2) made by each participant; probability matching should result in 7 such predictions and maximizing 10.

What evidence is there to suggest that, of the two strategies commonly used in this task, matching is the more intuitive, "dumb" response? As alluded to above, one benchmark involves asking the participants themselves which strategy is superior.

After they had played the cups game, but before the cups were turned over and they learned their payoff, participants were presented with two alternative strategies that could have been used in the game, one ascribed to a character named "Mike," who had guessed green for all 10 pairs of cups, and the other to "John," who had guessed green for 7 pairs and red for the remaining 3 pairs. Thus, Mike and John's strategies corresponded to maximizing and matching, respectively. When asked whose strategy most closely resembled their own, participants tended to answer in accord with the choices they had just made during the cups task. But, when asked whose strategy would be expected to earn more money, and whose strategy they would use if they were to play again, a substantial proportion (over 40%) of participants who had matched on the cups task nonetheless selected maximizing as the superior strategy. A substantially larger proportion of participants was able to identify maximizing as the superior strategy in a direct comparison to matching than had actually engaged in strict maximizing on the cups task itself. A similar finding was reported by Koehler and James (2009). In effect, many participants who engaged in probability matching in the binary prediction task later acknowledged that they would have been better off using a maximizing strategy instead.

We take this as evidence that many participants themselves would categorize probability matching as a "dumb" or inferior strategy. In fact, it is possible that the results above underestimate the proportion of participants who agreed with this categorization upon being presented with the direct comparison of the matching and maximizing strategies, as some participants who had engaged in matching on the cups task might have been reluctant to acknowledge their mistake to the experimenter. On the other hand, it is also possible that participants' responses to the strategy questions, such as the one asking which strategy they would use if they were to play again, do not accurately reflect what they would actually do if given such an opportunity. In a second study, Koehler and James (2010) provided a more direct test by presenting a strategy comparison question to some participants before they completed the binary prediction (cups) task. That is, these participants were presented with the two strategies that could be used on the cups task (7 green and 3 red, or 10 green and 0 red) and asked which would be expected to earn more money, prior to completing the cups task for themselves. Compared to a control group that did not first complete the strategy question, those who had compared the two strategies directly in terms of expected earnings were more likely to engage in strict maximizing and less likely to engage in strict matching on the subsequent cups task.

The results of Koehler and James (2010) suggest that one reason why people engage in probability matching rather than maximizing, at least on the cups task, is that matching is a more highly available strategy than maximizing: When confronted with the prediction task, matching tends to spring to mind more readily as a possible response than does maximizing. When the two strategies are equated for availability, as in the strategy comparison questions that describe both strategies, the relative appeal of matching diminishes. As researchers, many of us find probability matching intriguing because we are puzzled by why people choose to match rather than to maximize. From the participant's perspective, however, it seems that they are not really choosing between matching and maximizing strategies, as maximizing may simply not have come to mind as an alternative strategy to matching. Instead, perhaps, if matching is the only strategy that comes readily to mind as a response to the prediction task, specific alternatives to matching (such as maximizing) need to be effortfully "unpacked" from the generic "do something other than matching" category of strategies. For many people, apparently, matching springs to mind as a response and it seems "good enough" that the effort that would have to be expended to generate alternative strategies does not seem worthwhile. (The effortfulness of generating the maximizing strategy is discussed further in the next section.)

We take the results above to suggest that the matching strategy comes to mind more quickly and effortlessly than does the maximizing strategy as a possible response to the binary prediction task. In other words, probability matching may be characterized as an intuitive response. The idea that probability matching reflects an intuitive response, that may or may not be overridden by effortful deliberation, has been suggested by other researchers as well (Kogler & Kuhberger, 2007; West & Stanovich, 2003). But this characterization begs the question of *why* probability matching might be generated by the intuitive system as a response to prediction tasks such as the cups task described above. We recently conducted some studies to answer this question, in which we attempted to connect probability matching to an important function of the intuitive system: the generation of expectations.

To illustrate, imagine a 10-sided die with 7 green sides and 3 red sides. The die is going to be rolled 10 times. For most people, we suggest, simply providing this description is enough to trigger the expectation that the die roll will come up green seven times and red three times.[2] Much of our mental machinery is dedicated to the generation of expectations and predictions (e.g., Bar, 2007). For instance, most adults can readily translate outcome probabilities (e.g., that there is a 70% chance of a green outcome on each roll) into expected frequencies over a repeated sequence (e.g., that in 10 rolls of the die, 7 green and 3 red outcomes are expected). In fact, previous research has documented people's tendency to expect, even for very short sequences, outcome relative frequencies to correspond to the long-run probabilities governing the random generation process (Kahneman & Tversky, 1972; Tversky & Kahneman, 1971). We characterize this expectation generation process as an operation of the intuitive system. As with other such operations, its adaptive benefits are obvious. In the case of the 10-sided die, for example, if participants were asked to predict how many greens and how many reds would be rolled, then this expectation is exactly what is called for by the task. Precisely because of its usefulness in many predictive tasks, we suggest that expectation generation is the type of operation one might assume would migrate, with practice or experience, to the

[2] Here, we focus on the case of "described" prediction tasks in which participants are informed from the outset of the relevant outcome probabilities. It seems plausible that people would generate similar expectations in the case in which outcome probabilities are learned through observation, but it should be noted that the research we review in the remainder of this section involved described prediction tasks only.

intuitive system, such that expectations of this sort can be generated quickly, effortlessly, and without prompting from the deliberative system (e.g., Kahneman & Klein, 2009).

For the binary prediction task, however, expectation generation may not be entirely helpful. In the 10-sided die example, for instance, if one's task is simply to predict, prior to each of the 10 rolls, whether it will come up green or red, the best course of action is to maximize. Maximizing, in turn, requires only that the more likely outcome be predicted on every roll. As long as green is the more likely outcome, from the perspective of maximizing its precise probability does not matter, nor does the expected frequency of its occurrence over the sequence (i.e., 7 out of 10 rolls). But suppose the intuitive system nonetheless generates expected frequencies, quickly and effortlessly, such that they come to mind whether or not they are needed for the particular task at hand. Kahneman and Frederick (2002) describe a process of *attribute substitution* in which a heuristic attribute is rapidly evaluated via operations of the intuitive system and then—due to lax monitoring of the deliberative system—is substituted for the evaluation of the target attribute that is the intended focus of judgment. We use the notion of attribute substitution to explain the intuitive appeal of probability matching: expected frequencies generated by the intuitive system (e.g., expect seven greens and three reds) are in turn used to guide selection of a congruent prediction strategy (e.g., predict seven greens and three reds).

An important feature of this account, which we have referred to as *expectation matching*, is its focus on expected outcome frequencies over a sequence of events. In the 10-sided die example, the expectation that is evoked regards the sequence of 10 rolls as a whole: Over that sequence, we expect to see seven green rolls and three red rolls. A testable prediction of the expectation matching account, then, is that manipulations that disrupt or block the generation of sequence-wide expectations should reduce the rate of probability-matching behavior.

We conducted three experiments to test this prediction (James & Koehler, 2011). Each involved a sequence of 10 outcomes in a binary prediction task, in which the probability of one outcome was always 70% and the other was always 30%, as in the 10-sided die example we have been using here. We reasoned that when a single event (or "game") was played 10 times, as in the example of repeatedly rolling the 10-sided die, people would readily generate a sequence-wide expectation (e.g., 7 greens and 3 reds), which in turn was expected to foster probability matching in the prediction task. By contrast, if the 10 events or games were more individuated or distinct from one another, we reasoned, people would be less prone to generate (or apply) a sequence-wide expectation over the diverse collection of events or games, even if the outcome probabilities were the same, and therefore would be less likely to engage in probability matching.

In our first experiment, a unique games condition was created by asking participants to make binary predictions of the outcomes of 10 different games, which involved such activities as drawing ping-pong balls from a bingo cage, spinning a wheel of fortune and rolling a 10-sided die. In each game, as participants were informed in advance of making their predictions, one of the outcomes had a 70% probability of occurring and the other had a 30% probability. In the repeated games condition, 1 of the 10 games from the unique games condition was randomly selected for each participant and presented to him or her 10 times. Participants faced mathematically equivalent prediction tasks in the two conditions, which differed only in the superficial features that individuated the games in the unique-games but not in the repeated-games condition. Nonetheless, the rates of matching and maximizing in the two tasks differed significantly, as predicted: Participants in the unique games condition were less likely than those in the repeated games condition to engage in strict matching (3% vs. 38% of participants), and more likely to engage in strict maximizing (70% vs. 44% of participants).

A similar result was obtained in a second experiment (James & Koehler, 2011, Experiment 2), in which a 10-roll die game either involved the same die (with red and green sides) rolled 10 times or 10 different dice with unique markings (triangles vs. squares, hearts vs. flowers, etc.) which were each rolled once. In both experiments, apparently, individuating the sequence of outcomes made it less likely that participants generated or applied a sequence-wide expectation in making their predictions, reducing the rate of probability matching as a consequence. A third experiment presented an identically described prediction sequence to all participants, but preceded it by a priming manipulation designed to focus attention either on the sequence as a whole or on the individual outcomes within the sequence. This was accomplished by asking participants, after the prediction task involving the 10-sided die had been described to them, either to indicate in how many of the 10 rolls they expected each outcome (global focus condition) or to indicate, on any individual roll of the die, which outcome was more likely (local focus condition). We assumed that the global focus condition encourages generation of a sequence-wide expectation and the local focus condition does not. As hypothesized on this assumption, participants in the global focus condition were more likely to match and less likely to maximize on the prediction task than were those in the local focus condition.

In summary, the characterization of probability matching as "dumb" depicts it as arising from a fast intuitive process that is not reliably overridden by subsequent deliberation. Two lines of evidence support this characterization. First, matching occurs less frequently when the alternative maximizing strategy is brought explicitly to participants' attention, consistent with the claim that matching comes to mind quickly and spontaneously while maximizing does not. Second, consistent with the idea that matching results from an intuitive process that generates sequence-wide expectations, manipulations that individuate the sequence or otherwise encourage a focus on single outcomes decrease the rate of matching and increase the rate of maximizing.

3. SMART MAXIMIZING

On an account that characterizes probability matching as "dumb" in the sense of being a mistake rooted in the operations of the intuitive system, maximizing must be characterized as "smart" in the sense of representing avoidance or correction of that mistake through operations of the deliberative system. What evidence is there that maximizing is "smart" in this sense?

One strand of evidence suggesting that maximizing requires effortful deliberation comes from a study by Shanks, Tunney, and McCarthy (2002). Their participants made binary predictions in a standard probability-learning task and were paid for each correct prediction. After every 50-trial block, participants received a summary of their proportion of correct predictions on the block, and also the proportion correct that could have been obtained using an "optimal strategy." The number of participants who engaged in strict maximizing, which was defined as predicting the higher probability outcome on at least 50 consecutive trials, was examined as a function of the number of prediction blocks completed. In their Experiment 1, which consisted of 300 trials, only 6 of 16 participants were categorized as having engaged in strict maximizing. In their Experiment 2, in which the number of trials was increased to 1800 trials, 8 out of 12 participants eventually engaged in strict maximizing. What is striking to us about this result is how difficult it is, apparently, for participants to generate and consistently use the maximizing strategy despite repeated suggestions that a better strategy than matching is available and provision of hundreds of trials to identify it.

A natural starting point for discussing "smart maximizing" is the observation that not everybody engages in probability matching in binary prediction. As illustrated in Fig. 3.2, the typical finding is that some people match while others maximize. What individual difference variables distinguish these two groups? Broadly speaking, probability matching is more likely to be overridden in favor of maximizing when the individual is willing to engage in deliberation (i.e., has the appropriate motivation or thinking disposition) and has mastery of the basic normative principles (e.g., the calculation of expected value) needed to identify maximizing as the superior strategy. The distinction between these two components of deliberation has recently moved to the forefront of theoretical development of the idea of distinct thinking systems (e.g., Evans & Stanovich, 2013; Stanovich, 2009), but for our purpose, we largely gloss over this distinction and focus broadly on variables related to either facet of deliberative ability. On the smart maximizing account, we would expect variables that measure the propensity to rely on deliberation, or the effectiveness of deliberation, to positively correlate with the use of maximizing in prediction.

West and Stanovich (2003) found, in three studies, that maximizers tended to score higher than probability matchers on a measure of cognitive ability (self-reported total scores on the SAT Reasoning Test). This result, which was replicated in two subsequent studies by Stanovich and West (2008), is consistent with the idea that maximizing is fostered by processes of deliberation that are executed more reliably and efficiently by those of greater cognitive ability. Interestingly, in their studies, West and Stanovich did not find an association between the tendency to maximize (vs. match) and the number of math or statistics courses the participants reported having taken. This result could be taken to suggest that it is general deliberative ability, rather than specific mathematical knowledge, that promotes the use of maximization over matching. West and Stanovich did find, however, that matchers and maximizers differed in some important respects with regard to their perceptions or beliefs about the probabilities governing the outcomes in the prediction task. Specifically, matchers were significantly more likely than maximizers to endorse the gambler's fallacy that a long streak of one outcome made the alternative outcome more likely, while maximizers were more likely to endorse the notion of serial independence of outcomes (see also Gal & Baron, 1996).

Cognitive ability, then, which might be thought of as a measure of a person's ability to engage in effective deliberation, is associated with the tendency to maximize rather than to match, supporting the notion of "smart maximizing." Another, possibly related measure concerns *cognitive reflection*, or a person's tendency to scrutinize rather than unreservedly accept their initial, intuitive response to a problem or decision. In a highly influential paper, Frederick (2005) developed a brief cognitive reflection test (CRT) that can be taken as a measure of individual differences in proneness toward cognitive reflection (vs. reliance on intuition). The CRT consists of three mathematical problems, each of which has an "intuitive" but incorrect answer that many people report comes readily to mind. Correct responding, therefore, requires the person to override or correct that initial intuitive response. Toplak, West, and Stanovich (2011) report that CRT scores are independently, and more strongly, predictive of scores on a battery of "heuristics and biases" tasks (which included two probability-matching tasks) than is a measure of cognitive ability. On an account in which probability matching is the intuitive response, which must be overridden via deliberation to arrive at a maximizing strategy instead, higher CRT scores would be expected to be associated with a tendency to maximize rather than to match.

In the study involving the cups task described previously, we subsequently administered the CRT to all participants. Figure 3.3 relates performance on the prediction task to CRT scores in Study 1 of Koehler and

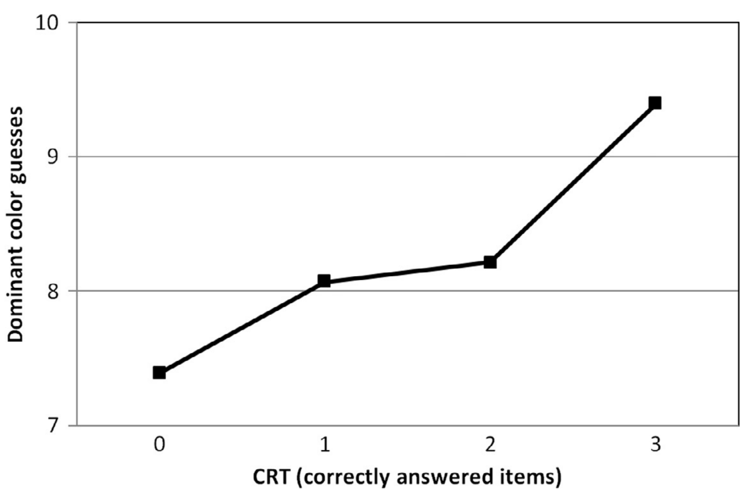

Figure 3.3 Mean number of predictions of the more probable outcome, out of a possible 10, on the cups task of Koehler and James (2010, Study 1), by CRT score.

James (2010). The mean number of times the dominant color (green in the example we have been using) was chosen is compared for those scoring 0, 1, 2, and 3 correct answers on the CRT. The figure reflects the substantial positive correlation between the two variables ($r = 0.40$): Those scoring higher on the CRT were more likely to maximize, and less likely to match, than those scoring lower on the CRT. Indeed, the mean number of green guesses for those scoring 3 out of 3 correct on the CRT is approaching that expected under strict maximizing, while the mean number of green guesses among those scoring 0 out of 3 on the CRT approaches that expected under strict matching. The relation between CRT score and maximizing remained statistically significant even after controlling for mathematical ability, as measured either by self-reported proficiency or by the number of math courses taken.

In one study involving the cups task (Koehler & James, 2010, Study 2), we also subsequently administered another task designed to measure proneness to reliance on intuition in decision making, namely, a variant of Epstein's jelly beans task (see, e.g., Denes-Raj & Epstein, 1994). Epstein and colleagues have depicted the jelly beans task as putting the intuitive and deliberative systems into conflict. Pacini and Epstein (1999) found that performance on the jelly beans task related to scores on the rational thinking component of the Rational–Experiential Inventory, a measure of individual differences in thinking dispositions. In our version of the task, participants were asked to consider two urns, one containing 1 gold ball and 9 white balls and the other containing 9 gold balls and 91 white balls. Participants were instructed to imagine that they are to make a single draw, at random, from one of the urns and that if a gold ball is drawn, they would 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 (a phenomenon commonly referred to as "ratio bias"). 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. A composite sum of these ratings correlated significantly with choices on the cups task, such that higher ratings in favor of the inferior urn (greater absolute number but smaller proportion of gold balls) were associated with a tendency to match rather than maximize on the cups task. Along with the results from the CRT, these findings support the idea that individual differences in the tendency to rely on (vs. override) an initial intuition are predictive of probability-matching behavior, as would be expected on a "smart maximizing" (or "dumb matching") account.

The "smart maximizing" account further predicts that experimental manipulations that foster deliberation ought to decrease the rate of probability matching in favor of maximizing. Surprisingly, few studies have tested this prediction directly. As is discussed further below, manipulations of cognitive load or working memory capacity could reasonably be expected to *reduce* the activity or involvement of the deliberative system. Therefore, on the smart maximizing account, they would be expected to increase the rate of matching. Unfortunately, interpretation of such studies is complicated—particularly when studied in the probability-learning paradigm—because such manipulations may also have effects on other mental operations such as those involved in detecting patterns or in monitoring and responding on the basis of observed outcome frequencies, and not just on the extent to which deliberation is used in selecting a prediction strategy. Here, we focus instead on manipulations intended to *increase* the general level of involvement of the deliberative system in overseeing and potentially correcting the output of the intuitive system.

The most direct evidence we know of comes from a study by Kogler and Kuhberger (2007), who studied probability matching in a task developed by Rubinstein (2002) that involves selecting 5 cards at random from a deck with known composition (36 green cards, 25 blue, 22 yellow, and 17 red), placing them in 5 envelopes, and asking participants to guess the color of the card contained in each envelope. Note that, even when selecting without replacement from the deck, the green card is the most likely outcome for each draw and therefore constitutes the maximizing response. Matching involves a representative selection of guesses that more closely reproduces the proportion of cards of each color in the deck. Previous work by the authors, and by Rubinstein (2002), showed that many people failed to maximize on this task and instead engaged in something more closely resembling a matching strategy. Kogler and Kuhberger (2007) compared responses from a control group, in which the task was likened to a lottery, to those assigned to a "corrective" condition in which the task was described as a statistical test designed to assess their level of statistical competence. Participants in the corrective condition were also advised to take their time in responding and to "carefully reconsider" their initial predictions before making a final response. The rate of maximizing was nearly three times higher among participants in the corrective condition (43%) than among those in the control condition (15%). Matching, by contrast, was more prevalent in the control (61%) than in the corrective (37%) condition. We take this result as supportive evidence for the claim that maximizing is a "smart" response that is fostered by enhanced deliberation, though of course other interpretations are also possible. In another strand of potentially supportive evidence, Fantino and Esfandiari (2002) found that telling participants that they would be asked, following the prediction task, to recommend a strategy to another participant—which arguably would motivate participants to think harder about the rationale for their predictions on the task—also increased the rate of maximizing. More recently, Taylor, Landy, and Ross (2012) found that providing a causal explanation for why one probabilistic outcome is more likely than another reduced the rate of probability matching in favor of a more optimal prediction strategy (though typically not strict maximizing), which could have resulted from the causal explanation triggering more extensive or rule-based deliberation.

In summary, the "smart maximizing" strategy depicts maximizing behavior as the product of a deliberative process that overrides the initial, intuitive tendency to engage in probability matching. Evidence supporting this characterization comes from correlational studies showing a positive association between maximizing and individual difference measures of cognitive ability and thinking disposition, and from experiments demonstrating increased maximizing as the result of manipulations designed to encourage greater deliberation.

4. SMART MATCHING

In the dual-systems account we have provided thus far, maximizing has been construed as resulting from deliberation and probability matching the result of a fast effortless heuristic, which we have labeled expectation matching. We have argued, therefore, that probability matching is essentially "dumb," but this view is not shared by all researchers who have studied the phenomenon. Indeed, probability-matching behavior has also been characterized as a sophisticated and adaptive response to an uncertainty about the true random nature of the binary prediction task (e.g., Gaissmaier & Schooler, 2008; Green et al., 2010). Evidence provided for this argument rests on two basic lines of research: (A) That probability matching tends to occur less in situations that tax cognitive resources and (B) that probability matching seems to be related to an effortful exploration of the environment (in this case, the sequence of outcomes observed in the binary prediction task). We will review both of these lines of evidence and evaluate how strongly they support the claim that probability matching is smart.

Return for a minute to our earlier argument that probability matching is intuitive. Typically, within a dual-systems framework, system 2 (deliberative) responses are considered to be more effortful or the result of more complex cognition and more easily disrupted by conditions that tax cognitive resources (e.g., Masicampo & Baumeister, 2008). By contrast, system 1 (intuitive) responses tend to be relatively effortless and less dependent on availability of cognitive resources. Given this, we could reasonably predict that probability matching should require less cognitive effort and resources. But research, reviewed below, has reached the opposite conclusion. This research suggests instead that probability matching is the result of effortful cognition and that maximizing occurs when cognitive resources are taxed.

For example, a number of studies have demonstrated that probability matching is associated with more effort or cognitive complexity than maximizing. Unturbe and Corominas (2007) and McMahon and Scheel (2010) both asked participants to report any strategies they used during the binary prediction task and found that the complexity of the patterns or rules reported by participants were inversely correlated with their tendency to choose the more probable option. In other words, probability matchers reported using more complex rules or strategies than did maximizers, a finding that appears to be at odds with the notion that probability matching is "dumb" and maximizing is "smart." Instead, it would seem to support the notion of "smart matching."

In addition, studies have also demonstrated that probability matching *decreases* under conditions of increasing cognitive load. (Recall that if probability matching is intuitive we would predict the opposite.) Wolford, Newman, Miller, and Wig (2004) demonstrated that giving participants a dual task that competed for left-hemisphere resources resulted in a decrease in probability matching. This result, however, was not replicated by Otto, Taylor, and Markman (2011) who failed to find any difference between rates of probability matching and maximizing under conditions of load versus no load. Nevertheless, even this result is potentially problematic for a dual-system account of probability matching (in which probability matching is "dumb"), because this account implies that probability matching ought to *increase* under load, which was not observed in Wolford et al. or in Otto et al.

As an alternative to manipulating working memory or cognitive load, researchers have recently attempted to reduce the engagement of the deliberative system through manipulations that decrease the supply of glucose available to the brain (e.g., Donohoe & Benton, 1999; Masicampo &

Baumeister, 2008). If probability matching is "dumb," we would predict that a lower supply of glucose, on which deliberation is claimed to be dependent, should lead to more matching, but in a study by McMahon and Scheel (2010), it was found that depleting glucose led to more maximizing and less probability matching. Again, this result could be taken to imply that probability matching is "smart" after all.

However, there are two ways to interpret the results of the studies presented above. One interpretation is that probability matching is *not* "dumb" (the product of effortless intuition), but "smart" (the product of effortful deliberation). While this is a viable interpretation, if adopted, we have to reconcile it with the evidence provided earlier that suggests that matching is "dumb." Another possibility is that probability matching is intuitively generated as a response strategy, but it is effortful to implement in the standard binary prediction task. That is, concluding that you should allocate outcome predictions in proportion to their probability may be intuitive, but actually implementing a matching strategy is relatively effortful (e.g., because it requires monitoring the relative frequency with which predictions are made over a series of trials). By contrast, once you have decided to maximize (even if arriving at that strategy was effortful), all you have to do to implement it is to predict the same outcome over and over again. Currently, we know of no research that directly tests these two possible interpretations of why the prevalence of probability matching decreases under load, so we cannot yet definitively answer whether or not the findings described above challenge the notion that probability matching is "dumb."

Some researchers do argue, however, for the "smart" interpretation of probability matching, suggesting specifically that it represents misapplication or overgeneralization of the usually adaptive tendency to seek patterns or predictability in outcomes observed over time. This account rests on the notion that people tend to have misconceptions and misperceptions of randomness (see Falk & Konold, 1997 for review). Proponents of the "smart" account of probability matching argue that participants do not believe that the sequence they are exposed to in the binary prediction task is truly random, and they attempt to outperform the optimal maximizing strategy by finding a predictable outcome pattern that can be exploited to achieve better predictive accuracy (Gaissmaier & Schooler, 2008; Peterson & Ulehla, 1965). By this account (henceforth called the pattern-search hypothesis), probability matching is seen as a by-product of pattern search, in the sense that, if there actually was a predictable outcome pattern, exploiting it would produce predictions that are made in proportion to the relevant outcome probabilities. Because pattern search seems to require effortful deliberation, any resulting probability matching could be viewed as "smart" in comparison to the relatively "dumb" maximizing strategy of simply predicting the more likely outcome on every trial.

In support of the pattern-search hypothesis, several researchers have demonstrated that increasing the perceived randomness of the sequence, either by emphasizing its randomness (Morse & Runquist, 1960) or by making it appear more random[3] by increasing the frequency of alternations between different outcomes (Wolford et al., 2004), leads to more maximizing behavior. They argue that this is because the increase in apparent randomness of the outcome sequence disrupts participants' initial assumption that the sequence is nonrandom, leading participants to abandon their pattern search.

Further support for the pattern-search account comes from probability matchers' superior ability to detect patterns when they do exist. Gaissmaier and Schooler (2008) divided a typical binary prediction task into two halves. In the first half, participants were presented with a truly random sequence. From these predictions, Gaissmaier and Schooler classified participants as either probability matchers or maximizers. In the second half of the experiment, participants continued the binary prediction task, but this time, there was a nonrandom pattern in the outcome sequence. Probability matchers were significantly more likely to detect and exploit this pattern (as measured by increased prediction accuracy) than were maximizers. Gaissmaier and Schooler's work provides evidence that a probability-matching strategy, hypothesized to be grounded in pattern search, can convey an advantage in situations where patterns exist and is in this sense smart. But their results do not provide evidence that pattern search necessarily causes probability matching.

In fact, we provide evidence in Koehler and James (2009) against such a causal relationship. We asked participants to complete a computer-based binary prediction task in which they were to guess the color of marbles drawn from a bag. The bag consisted of a mix of 30 green marbles and 10 red marbles.[4] Both groups participated in a learning phase in which they

[3] Ironically, this actually makes it less random, but the important aspect is how it appears to participants.

[4] For ease of discussion, we will always refer to green as being the more probable colour. In fact, this was counterbalanced so that for half of the participants, the contingencies were reversed and there were 30 red marbles and 10 green marbles.

were told that 40 marbles had been drawn randomly with replacement from the bag. Participants learned about the results of this set of draws in one of two ways: they were either told that result was 30 green marbles and 10 red marbles (the aggregate learning phase), or each marble was displayed serially as it was drawn (the serial learning phase). Note that those in the serial learning phase had the opportunity to search for patterns in the sequence of draw outcomes, but those in the aggregate learning phase did not.

After completing the learning phase, participants advanced to the testing phase in which they were asked to predict the colors of 20 more marbles that would be drawn from the bag (again with replacement). For each correct prediction participants could earn $0.50. Participants either made their guesses one at a time with no feedback on their accuracy (the serial testing phase), or they indicated in aggregate how many times they intended to guess green and how many times they intended to guess red (the aggregate testing phase). Again, those in the serial testing phase had the opportunity to include pattern information in their response if they chose to, while those in the aggregate condition did not.

If probability matching is the result of a pattern search, then we would expect those given the opportunity to observe (serial learning) and exploit (serial testing) pattern information to exhibit more probability matching. In fact, we found no difference between those with access to pattern information at learning or test and those without it in terms of their prediction strategies. All conditions showed an equal and high degree of probability-matching behavior.

The results of Koehler and James (2009) suggest that, while pattern search may indeed be associated with probability matching as Gaissmaier and Schooler demonstrated, it does not (at least explicitly) appear to cause participants to probability match. More generally, it is not clear why searching for patterns itself would lead participants to probability match, assuming that they have not been successful in finding a pattern. Why couldn't a participant implement a maximizing strategy (once the outcome with the higher probability has been identified) while also keeping an eye out for patterns that, if discovered, could be exploited for greater predictive accuracy? In fact, as maximizing seems relatively effortless to implement, it would seem like the ideal candidate strategy to use while devoting the bulk of one's cognitive resources to searching for patterns.

Given that pattern information is not necessary for probability matching, and that searching for patterns does not actually require the use of a probability-matching strategy, it seems premature to conclude that pattern search causes probability matching. Indeed, the alternative interpretation (that probability matching leads to pattern search) seems equally plausible. One could argue that, having decided to probability match (arising from expectation matching by the intuitive system), participants realize that they are only part way to perfectly predicting the outcome. The vital piece of information they are missing is how to order their predictions. To fill the gap, participants begin searching for patterns. This is, of course, an ad hoc explanation that requires testing. But it does illustrate that we should be cautious in concluding that pattern search causes probability matching.

Along similar lines, if we assume that probability matching causes pattern search (instead of the other way around), we can also offer an alternative explanation of the work investigating the role of perceived randomness in dual choice tasks. Recall that when outcomes were constructed so that they appeared more random (i.e., by introducing a higher rate of alternation), it led to more maximizing behavior. Currently, this is the strongest evidence for a pattern search account (even though it is not a direct test of it), largely because it establishes causation. Assuming that pattern search does not cause probability matching requires some alternative explanation of this data. For illustrative purposes, imagine that you are engaged in a search for patterns because you have decided to probability match. As it becomes obvious through search that you will not find any pattern information, you may revisit (deliberate on) your initial strategy choice, which may cause you to switch to maximizing. Thus, if the sequence appears to have no patterns, you abandon your pattern search sooner. This interpretation is not all that different from that of the original researchers investigating perceived randomness and matching, except that it assumes that abandoning pattern search leads participants to deliberate. Indeed, in the original research, it is not clear why or how participants are said to arrive at a maximizing strategy after they abandon their matching approach. Apparently, it is assumed that maximizing is what participants would have done if they had not chosen to match instead, but this assumption remains unelaborated in the original work.

In summary, is there any conclusive evidence to suggest that probability matching is smart? The short answer, we argue, is not yet. Manipulations designed to tax cognitive resources have not been shown to increase probability-matching behavior, which we would definitely expect on the "dumb matching" account. But more work needs to be done to determine whether the results from concurrent load and depletion manipulations come about because probability matching is actually "smart," or simply because it takes more effort to implement. Those arguing that probability matching represents an overgeneralization of a usually intelligent search for patterns in outcome sequences might have a good case, but they still need to establish a solid causal relationship. However "smart" the mechanism that produces it, probability matching that arises from a fruitless search for patterns in a truly random sequence is at least "dumb" in the sense of being suboptimal and costly in that particular setting.

5. DUMB MAXIMIZING

Arguments that probability matching is the result of a deliberative search for patterns have largely neglected to discuss where maximizing comes from. It remains unclear why a strategy that is the result of a calculated analysis, according to those who think maximizing is smart, is also the default that people revert to when they are too taxed to search for patterns. In this case, it seems that maximizing is taking on the role of the dumb strategy, but this begs the question, is there a "dumb" mechanism by which maximizing can be produced?

There is, in fact, substantial evidence, suggesting that maximizing can arise from system 1 mechanisms. Although little work has investigated the issue directly, many studies have demonstrated maximizing (rather than matching) in populations that are not prone to, or efficient at carrying out, extensive deliberation. For example, children have been found to maximize when they are very young and do so even when they cannot report explicitly which event is more likely (Derks & Paclisanu, 1967). Furthermore, maximizing has been reported in a number of nonhuman species (Parducci & Polt, 1958; Wilson, 1960), who presumably do not possess the deliberative ability required to identify maximizing as the superior strategy in terms of expected value. Aside from these findings, there is also the work presented earlier demonstrating that when under cognitive load (Wolford et al., 2004) or when deprived of glucose (McMahon & Scheel, 2010), maximizing behavior also increases. Under these situations, we would expect intuitive responding, so this lends further support to the notion that there may be a "dumb" variant of maximizing.

What is the mechanism behind this intuitive, dumb maximizing? As discussed earlier, it could arise from the relative simplicity of implementing the maximizing strategy, but it could also be the result of basic operant conditioning. In situations with outcome (or reward) feedback, predicting the more probable outcome will be rewarded more frequently than will predicting the less likely outcome. Given this asymmetry of reward, operant conditioning eventually should produce maximizing behavior. Indeed, it was the assumption of the early animal literature investigating binary prediction problems that maximizing should be the default (Parducci & Polt, 1958; Wilson, 1960).

Despite the clear opportunity of a role for operant conditioning in producing maximizing behavior, little attention has been paid to it in recent work on human probability matching. McMahon and Scheel allude to it in arguing that glucose depletion should lead to maximizing because "predicting the most frequent outcome produces the highest rate of reinforcement" (McMahon & Scheel, 2010, p. 450). The general assumption in the recent literature, however, is that maximizing should prevail because it is the (deliberatively) rational solution to the binary prediction task. From a dual-system perspective, operant conditioning represents an intuitive operation of system 1, as opposed to deliberative strategy generation and comparison processes arising from system 2. But operant conditioning is quite distinct from the intuitive process (expectation matching) we argued gave rise to probability matching. Although they are both intuitive, the processes differ in two key ways: (a) the effects of operant conditioning are bottom-up (data driven) and potentially unavailable to awareness, whereas expectation matching produces a top-down (theory driven) solution that is highly available to awareness; and (b) operant conditioning can foster maximizing only in situations where outcome or reward feedback is present, while this limitation does not pertain to expectation matching. Each of these points has important ramifications for the predictions of a "dumb maximizing" account and also for how we can understand dumb maximizing within the context of dual systems, so we will discuss each of these points in turn.

To our knowledge, little research currently exists on how or when operant conditioning might influence strategy in a binary prediction task. The only available evidence that we know of comes in the form of response times. Participants are generally slower to choose the less probable option than they are to choose the more probable option in the binary prediction task (Otto et al., 2011; Unturbe & Corominas, 2007).[5] One interpretation of this finding is that participants are slower when making a prediction of the

[5] One initial problem with this finding is that we would expect participants to be faster if they are choosing the same response twice in a row, which they would do more often with the more probable option. More careful analysis is needed to ensure that reaction time is slower for low probability outcome predictions independent of whether it is a repeated choice or not.

lower probability outcome because operant conditioning is pushing them to choose the more probable option, and the conflict takes time to resolve.

This interpretation suggests the two strategies (top-down matching and bottom-up maximizing) can jointly influence execution of participants' predictions (Newell, Koehler, James, Rakow, & van Ravenzwaaij, 2013), but this interaction has not been extensively studied. For instance, it is not clear to what extent operant conditioning impacts strategy choice, or whether its operations are available to conscious awareness. To illustrate, consider a participant who has chosen, via a top-down strategy selection process (e.g., guided by expectation matching), to probability match. How will operant conditioning impact this participant's choices? It is possible that the impact will be small or inconsequential in comparison to their explicitly adopted strategy to probability match and will therefore have little effect. Alternatively, if the effects of operant conditioning are available to awareness, it might cause the participant to deliberate and thereby increase the incidence of top-down maximizing. Finally, operant conditioning could have an unconscious and subtle influence on choices leading the participant to blend maximizing and matching without being explicitly aware of taking that approach.[6]

Unfortunately, the current literature does not provide much basis for differentiating between these alternative accounts. If operant conditioning is having some effect on choices, we would expect to see more maximizing over time, in tasks that involve many trials and outcome or reward feedback. There is substantial evidence that the rate of maximizing increases over trials (Bereby-Meyer & Erev, 1998; Edwards, 1961), but there has been little controlled work done to investigate whether operant conditioning is the cause of this trend. In addition, the data do not allow determination of whether the effects of operant conditioning might be conscious or unconscious. If, for example, this trending is the result of operant conditioning encouraging deliberation, we might expect some individuals to switch abruptly from another strategy (e.g., matching) to strict maximizing. But, if operant conditioning is providing an unconscious nudge, the change may be more gradual and incomplete; that is, it may never reach strict maximizing, but might instead exhibit what has been referred to as overmatching (Friedman & Massaro, 1998; see also Vulkan's, 2000 review). When looking at group

[6] This list of potential influences of operant conditioning is intended to be illustrative rather than exhaustive.

means, however, as is typically reported in the literature, it is difficult to distinguish these two possibilities.

It has not been established that overmatching is the result of operant conditioning (let alone whether or not the effect of that conditioning is conscious, unconscious or both), but there are a variety of ways to test whether effects of operant conditioning on binary prediction are (a) conscious and (b) contributing to overmatching. If operant conditioning fosters overmatching, for example, we should see overmatching occurring more frequently in situations with feedback and many trials. It should be a less common finding in studies with few trials, no feedback, and "described" contingencies (i.e., in which the relevant outcome probabilities are explicitly provided to participants rather than having to be estimated from trial-by-trial observation of outcomes). If operant conditioning is leading to probability matching unconsciously then, when queried, overmatchers should still report that probability matching is the optimal strategy. Such a finding would provide stronger evidence that a strategy approaching maximizing (such as overmatching) need not be coupled with explicit endorsement of a maximizing strategy. In fact, it could coexist with the actual intent to probability match! We hope to investigate this possibility in future research.

The second important feature of "dumb maximizing" worth discussion is that it can only operate in situations that include feedback (or, more specifically, administration of reward). This feature suggests some promising ways to test, in future research, whether operant conditioning is at work in dumb maximizing. For example, by varying whether or not participants receive feedback, we can also vary whether or not it is possible for operant conditioning to influence responses. Thus, manipulations that encourage reliance on system 1 operations (such as the cognitive load and glucose manipulations discussed earlier) should only lead to more maximizing behavior in situations with feedback, as under those circumstances greater reliance on processes of operant conditioning should foster maximizing. If feedback is not provided, we would expect these manipulations to increase probability matching (assuming that matching and maximizing are equated on implementation effort), as the only remaining intuition, when operant conditioning is not at work, is the expectation generation process that we have argued produces "dumb" matching.

It is also worth noting that feedback is only useful for producing operant conditioning that will influence binary prediction if it serves as a reward or punishment that is contingent on predictions. For this to be the case, reward administration must follow a prediction of some sort. To use the die problem, we have referenced throughout this chapter as an example, if you predicted green and the die comes up green, you receive a reward and that response is reinforced. If the die comes up red, you fail to receive a reward and that response is negatively reinforced. But if you made no guess, there is no possibility of reward or of reinforcement. Indeed, research shows that feedback improves performance on the binary prediction task, but only if it follows a prediction made by the participant. Observation-only trials do not improve performance (Newell & Rakow, 2007; Tversky & Edwards, 1966), consistent with the operant-conditioning account.

One final piece of evidence supporting the notion that operant conditioning encourages maximizing behavior even when probability matching is the explicitly selected strategy comes from work by Newell and Rakow (2007). In their version of the binary prediction task, outcome probabilities were fully described to participants before they made any guesses. Following this description, participants made a series of predictions with or without feedback. Newell and Rakow found that performance drifted toward maximizing over time to a significantly greater extent in the feedback condition than in the no-feedback condition. This result can be seen as something of a puzzle as the explicit system already had all the necessary information (i.e., the relevant outcome probabilities) to make an optimal choice before any predictions were made or any outcomes observed. From a rational choice perspective, the information provided by feedback was completely extraneous. One way of making sense of these data is to argue that the trend toward maximizing was brought on by operant conditioning, which is necessarily inactive in conditions without feedback, but in conditions with feedback, it is able to slowly push choices toward maximizing. While many interpretations of this finding are possible, it is consistent with the notion that feedback encourages optimal responding through operant conditioning.

In summary, it is highly plausible that maximizing may arise, over the course of trial-by-trial experience, as the consequence of bottom-up processes (e.g., through mechanisms of operant conditioning) that might be characterized as relatively "dumb" (e.g., in comparison with "smart" deliberative processes that lead to top-down identification of maximizing as the superior predictive strategy). Some supportive evidence for dumb maximizing comes from studies demonstrating (a) the importance of active prediction in increasing maximizing rates over trials and (b) an increase in maximizing rates over trials with feedback even in fully described prediction tasks, as well as from response time data.

6. CONCLUSION

An overarching theme of this chapter is that probability matching, and maximizing, behavior should not necessarily be taken as the product of a single process. Instead, there may be different processes, some relatively "smart" and others relatively "dumb," that give rise to either type of behavior depending on the circumstances in which it is observed. We are not the first to note that there might be different variants of probability matching (see, e.g., Gaissmaier & Schooler, 2008; Otto et al., 2011). This chapter highlights the possibility that there might be more than one variant of maximizing, as well.

We find the dual-system approach to be helpful in organizing discussion of variants of probability matching and maximizing in terms of the mental operations that produce them. The dual-system approach, as it has been applied to date to the phenomenon of probability matching, has largely drawn attention to the "dumb" (intuitive, fast, effortless) variant of matching and to the "smart" (effortful, slow, deliberative) variant of maximizing, both in our own work and in that of other researchers (Kogler & Kuhberger, 2007; West & Stanovich, 2003). Our goal in this chapter was to expand the dual-system approach to encompass the complementary possibilities of "smart" matching and "dumb" maximizing. Without a more complete picture, we are left with the riddle of why children and nonhuman animals sometimes conform more closely in their predictions and decisions than do otherwise more sophisticated adult humans to the prevailing model of rational choice.

ACKNOWLEDGMENTS

Preparation of this chapter was supported by a grant from the Natural Sciences and Engineering Research Council of Canada to the first author. We are grateful to Ben Newell for helpful comments.

REFERENCES

Bar, M. (2007). The proactive brain: Using analogies and associations to generate predictions. *Trends in Cognitive Sciences*, *11*(7), 280–289.

Bereby-Meyer, Y., & Erev, I. (1998). On learning to become a successful loser: A comparison of alternative abstractions of learning processes in the loss domain. *Journal of Mathematical Psychology*, *42*(2–3), 266–286.

Denes-Raj, V., & Epstein, S. (1994). Conflict between intuitive and rational processing: When people behave against their better judgment. *Journal of Personality and Social Psychology*, *66*(5), 819–829.

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

Donohoe, R. T., & Benton, D. (1999). Cognitive functioning is susceptible to the level of blood glucose. *Psychopharmacology*, *145*(4), 378–385.

Edwards, W. (1961). Probability learning in 1000 trials. *Journal of Experimental Psychology*, *62*(4), 385–394.

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

Evans, J. S. B. T., & Stanovich, K. E. (2013). Dual-process theories of higher cognition: Advancing the debate. *Perspectives on Psychological Science*, *8*(3), 223–241.

Falk, R., & Konold, C. (1997). Making sense of randomness: Implicit encoding as a basis for judgment. *Psychological Review*, *104*(2), 301–318.

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.

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

Friedman, D., & Massaro, D. W. (1998). Understanding variability in binary and continuous choice. *Psychonomic Bulletin & Review*, *5*(3), 370–389.

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*(1), 81–98.

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

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), 1–5.

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

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), 707–711.

Kahneman, D. (2011). *Thinking, fast and slow*. New York, NY: Farrar, Straus and Giroux.

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). NewYork, NY: Cambridge University Press.

Kahneman, D., & Klein, G. (2009). Conditions for intuitive expertise: A failure to disagree. *American Psychologist*, *64*(6), 515–526.

Kahneman, D., & Tversky, A. (1972). Subjective probability: A judgment of representativeness. *Cognitive Psychology*, *3*(3), 430–454.

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

Koehler, D. J., & James, G. (2010). Probability matching and strategy availability. *Memory & Cognition*, *38*(6), 667–676.

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

Masicampo, E. J., & Baumeister, R. F. (2008). Toward a physiology of dual-process reasoning and judgment: Lemonade, willpower, and expensive rule-based analysis. *Psychological Science*, *19*(3), 255–260.

McMahon, A. J., & Scheel, M. H. (2010). Glucose promotes controlled processing: Matching, maximizing, and root beer. *Judgment and Decision Making*, *5*(6), 450–457.

Morse, E. B., & Runquist, W. N. (1960). Probability-matching with an unscheduled random sequence. *American Journal of Psychology*, *73*, 603–607.

Newell, B. R., Koehler, D. J., James, G., Rakow, T., & van Ravenzwaaij, D. (2013). Probability matching in risky choice: The interplay of feedback and strategy availability. *Memory & Cognition*, *41*, 329–338.

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

Otto, A. R., Taylor, E. G., & Markman, A. B. (2011). There are at least two kinds of probability matching: Evidence from a secondary task. *Cognition*, *118*(2), 274–279.

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 and Social Psychology*, *76*(6), 972–987.

Parducci, A., & Polt, J. (1958). Correction vs. noncorrection with changing reinforcement schedules. *Journal of Comparative and Physiological Psychology*, *51*(4), 492–495.

Peterson, C. R., & Ulehla, Z. J. (1965). Sequential patterns and maximizing. *Journal of Experimental Psychology*, *69*(1), 1–4.

Rubinstein, A. (2002). Irrational diversification in multiple decision problems. *European Economic Review*, *46*, 1369–1378.

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), 233–250.

Sloman, S. A. (1996). The empirical case for two systems of reasoning. *Psychological Bulletin*, *119*(1), 3–22.

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

Stanovich, K. E., & West, R. F. (2000). Individual differences in reasoning: Implications for the rationality debate? *Behavioral and Brain Sciences*, *23*(5), 645–665.

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), 672–695.

Taylor, E. G., Landy, D. H., & Ross, B. H. (2012). The effect of explanation in simple binary decision tasks. *The Quarterly Journal of Experimental Psychology*, *65*(7), 1361–1375.

Toplak, M. E., West, R. F., & Stanovich, K. E. (2011). The cognitive reflection test as a predictor of performance on heuristics-and-biases tasks. *Memory & Cognition*, *39*(7), 1275–1289.

Tversky, A., & Edwards, W. (1966). Information versus reward in binary choices. *Journal of Experimental Psychology*, *71*(5), 680–683.

Tversky, A., & Kahneman, D. (1971). Belief in the law of small numbers. *Psychological Bulletin*, *76*(2), 105–110.

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

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

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., Jr. (1960). Supplementary report: Two-choice behavior of monkeys. *Journal of Experimental Psychology*, *59*(3), 207–208.

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), 221–228.