Striving for perfection and falling short: The influence of goals on probability matching

Jie Gao · James E. Corter

Published online: 10 January 2015
© Psychonomic Society, Inc. 2015

Abstract Probability matching in sequential prediction tasks is argued to occur because participants implicitly adopt the unrealistic goal of perfect prediction of sequences. Biases in the understanding of randomness then lead them to generate mixed rather than pure sequences of predictions in attempting to achieve this goal. In Study 1, $N = 350$ participants predicted 100 trials of a binary-outcome event. Two factors were manipulated: probability bias (the outcomes were equiprobable or distributed with a 75 %–25 % bias), and goal type—namely, whether single-trial predictions or the perfect prediction of four-trial sequences was emphasized and rewarded. As we hypothesized, predicting sequences led to more probability-matching behavior than did predicting single trials, for both the bias and no-bias conditions. In Study 1B, we added a control condition to distinguish the effects of the grouped presentation of trials from the effects of sequence-level perfect-prediction rewards. The results supported goal type rather than presentation format as the cause of the Study 1 differences in matching between the sequence and single-trial conditions. In Study 2, all participants ($N = 300$) predicted the outcomes for five-trial sequences, but with different goal levels being rewarded: 60 %, 80 %, or 100 % correct predictions. The 100 % goal resulted in the most probability matching, as hypothesized. Paradoxically, using the inferior strategy of probability matching may be triggered by adopting an unrealistic perfect-prediction goal.

Keywords Decision making · Probability matching · Prediction · Goals · Heuristics

J. Gao · J. E. Corter (✉)
Department of Human Development, Teachers College, Columbia University, 525 W. 120th Street, New York, NY 10027, USA
e-mail: jec34@tc.columbia.edu

Present Address:
J. Gao
Educational Testing Service, Princeton, NJ, USA

Probability matching refers to people's tendency, when faced with the task of repeatedly trying to predict which of two binary outcomes will occur, to match their prediction probabilities for the two outcomes to the observed outcome probabilities (e.g., Edwards, 1956, 1961; Herrnstein, 1961; Voss, Thompson, & Keegan, 1959; see Myers, 1976, or Vulkan, 2000, for reviews). For example, when asked to repeatedly predict the outcome of rolling a die with four green and two red sides, many people allocate about two-thirds of their color predictions to green and one-third to red. This probability-matching prediction strategy is suboptimal; the optimal strategy, termed *maximizing*, is to choose the outcome with the higher probability in every trial in the prediction task. People's use of probability-matching strategies has been a long-standing conundrum in decision-making research, because this behavior violates normative economic models assuming that people tend to make choices that maximize expected utility.

Theoretical accounts of probability matching

People exhibit probability matching in at least two types of situations involving repeated predictions. The first type of task, termed “probability learning,” is one in which the probabilities of the binary outcomes are initially unknown (e.g., Edwards, 1961; Estes, 1964).

In probability learning, simple reinforcement learning models (e.g., Bush & Mosteller, 1955; Thorndike, 1898) have been proposed to account for the emergence of probability matching and of other superior patterns of response (e.g., Herrnstein, 1961; Herrnstein & Loveland, 1975).

But probability matching is also observed in prediction tasks in which the relevant outcome probabilities are known from the outset (e.g., Fantino & Esfandiari, 2002; Koehler & James, 2009, 2010; Newell, Koehler, James, Rakow, & van

Ravenzwaaij, 2013; Newell & Rakow, 2007). In these tasks probability matching cannot easily be ascribed to reinforcement learning among responses, yet it is still a robust phenomenon. Theories that have been advanced to account for probability matching in these situations include simple “System 1” processes, that might collectively be termed “underthinking” (e.g., Koehler & James, 2010; West & Stanovich, 2003), for example the simple heuristic decision strategy known as “win–stay/lose–shift” (Gaissmaier & Schooler, 2008; Nowak & Sigmund, 1993; Otto, Taylor, & Markman, 2011). Another theoretical explanation of probability matching is that it arises from a process of *expectation matching* (James & Koehler, 2011; Koehler & James, 2009; Kogler & Kuhberger, 2007; West & Stanovich, 2003). In expectation matching, “the individual stochastically and independently generates predictions in accordance with this historical assessment of outcome probabilities” (Otto et al., 2011, p. 274). Other researchers (e.g., Gaissmaier & Schooler, 2008; Unturbe & Corominas, 2007; Wolford, Newman, Miller, & Wig, 2004) have demonstrated that probability matching in description-based tasks can arise from more complex “System 2” cognitive processes (erroneously applied), thus constituting a form of bias due to “overthinking”—for example, explicit mental search for patterns in sequences (e.g., Gaissmaier & Schooler, 2008; Wolford, Miller, & Gazzaniga, 2000; Yellott, 1969).

Settling the question of why probability matching occurs may be especially difficult because probability-matching behavior might be produced by more than one type of mechanism. For example, sensitivity to patterns might result from a reaction to the perceptual salience of “runs” of a single outcome (Vitz, 1964; Wagenaar, 1970), or through intuitive cognitive biases such as imagined sequential dependencies. In fact, some researchers (e.g., Gaissmaier & Schooler, 2008; Otto et al., 2011) have argued that probability matching can be produced either by low-level “System 1” mechanisms (e.g., reinforcement learning, perceptual biases, or the use of simple decision heuristics), or by higher-level “System 2” processes (e.g., pattern search, statistical reasoning, or use of explicit decision strategies). This idea echoes evidence that *maximizing* behavior, too, can arise in very different settings, apparently due to different mechanisms (Newell et al., 2013): It is observed both in animals, presumably as a result of association learning processes (Vulkan, 2000), and in relatively high-ability human participants (West & Stanovich, 2003), apparently as a result of explicit higher-level reasoning processes (e.g., Gaissmaier & Schooler, 2008; Gal & Baron, 1996).

The effect of goals on matching and maximizing

We believe that taking account of an individual’s goals in decision making (e.g., Fishbach & Ferguson, 2007; Krantz

& Kunreuther, 2007; Kruglanski et al., 2002) can help explain when and why probability matching occurs. Specifically, we hypothesize that probability matching is more likely to occur when the participant implicitly adopts the high-performance, low-probability goal of *perfectly* predicting sequences of trials. Obviously, such a perfect-prediction goal becomes increasingly unrealistic as the length of the prediction string grows. Then, participants behave as if they believe that perfect prediction of such a sequence is most likely to be achieved by generating appropriate mixed strings of the two possible outcome (i.e., with expectation matching).

The idea that striving for perfect prediction of sequences might exacerbate probability matching is not new. Estes (1964) speculated that particular tasks might tend to imply a perfect-prediction goal and that this might lead to probability matching, though he did not elaborate on why this might be so. Arkes, Dawes, and Christensen (1986) suggested that high motivation, leading to a desire to “beat the rule,” may be a factor in participants’ failure to maximize. They also found that self-assessed expertise was associated with a *lower* rate of maximizing, implicating (unrealistic) overconfidence by the “experts” as a factor in probability matching. Fantino and Esfandiari (2002) directly tested the idea that probability matching might arise from adoption of a perfect prediction goal. In their study, half of the participants were instructed that they should expect to do no better than 75 % correct (the probability of the more likely outcome). This cue explicitly discouraged adoption of a perfect-prediction goal, thus it might be expected to boost maximizing. Overall, the main effect of this cue was not significant, but in certain conditions (i.e., when participants were informed in advance of the relevant outcome probabilities, and when the participant was asked to provide strategy hints to another student), the manipulation had a nontrivial effect in the expected direction. The present study was an attempt to provide a more conclusive test of this hypothesis, using different manipulations and with a large sample size to ensure sufficient power.

It may be helpful to speculate why adopting a perfect prediction goal might encourage probability-matching behavior. One possible mechanism is that in trying to achieve perfect performance, people will be more prone to generate a sequence that they believe is *representative* of a random binary sequence, because intuitively (but incorrectly), people believe that a string that “looks right” in the sense of having an appropriate mix of outcomes has a higher probability of occurring (Bar-Hillel, 1982; Bar-Hillel & Wagenaar, 1991; Kahneman & Tversky, 1972). This bias toward producing representative strings of predictions thus leads participants to match the frequency of the binary alternatives in their prediction strings to the observed or stipulated base rates, as has been described by expectation-matching accounts (e.g., Gal & Baron, 1996; James & Koehler, 2011; Kogler & Kuhberger, 2007).

In contrast, a maximizing strategy generates a sequence consisting solely of the more likely outcome; thus, it does not seem at all “representative” of a random sequence. The paradox here is that to the naïve participant, the maximizing strategy *seems* to have a lower probability of predicting perfectly than does the probability-matching strategy (although actually it is higher). For example, in a sequence of ten trials, in which the probability of event A in each trial is 80 %, the maximizing strategy response will result in a 10.7 % chance of perfect prediction [i.e., the binomial probability of ten successes in ten trials, given that the single-trial probability of a correct guess is $(1)(.8) = .8$], but a probability-matching strategy has only a 2.1 % chance of predicting perfectly [the same binomial probability, but with the single-trial probability of a correct guess being equal to $(.8)(.8) + (.2)(.2) = .68$]. Adopting a more realistic goal (lower than perfect performance) might prevent people from being overly influenced by a representativeness bias, making it easier for them to discover or consider the optimal maximizing strategy. The studies below were designed to test this idea.

Study 1

In Study 1, we investigated whether presenting the binary prediction task to the participant as involving the prediction of *sequences* rather than predicting repeated individual trials, and rewarding perfect performance for sequences, would tend to elicit more probability matching. Thus, in Study 1 we manipulated two task factors. One factor was the nature of the rewarded performance goal (perfect prediction of a *sequence* of trials or correct prediction of *individual trials*), and the other factor was the probability “bias” of the events, with the probability of the more likely outcome being set either at 75 % (bias) or 50 % (no bias). Previous studies of probability matching have not included a condition in which the binary outcomes were equally likely, because in this condition probability matching is not inferior to maximizing—any strategy produces 50 % correct responses on each trial. We included this condition to assess generalizability—specifically, to check whether rewarding perfect prediction of sequences rather than the correct prediction for single outcomes could affect the prevalence of matching behavior for “nonbiased” (i.e., equiprobable) outcomes—the sort of prediction task featured in the literature on representativeness (e.g., Bar-Hillel, 1982; Kahneman & Tversky, 1972).

Method

Participants The experiment was conducted via Amazon’s Mechanical Turk (AMT) worker marketplace. We restricted recruitment to US participants only, and only to those with a

task approval rate higher than or equal to 95 %. We took care that each participant performed the task only once. The total number of participants tested was 350. The participants were relatively diverse: Their mean age was 30.9 years, and ranged from 17 to 64. Participants were 46 % male and 54 % female; 81.7 % were English speakers, and the other participants’ native languages included Arabic, Chinese, French, and other languages.

Procedure The interface program was developed in the HTML language with embedded JavaScript. Upon starting the task, participants were randomly assigned to one of the four conditions (sequence_bias, sequence_no-bias, single-trial_bias, and single-trial_no-bias), resulting in *N*s of 87, 87, 88, and 88, respectively. They were then shown the main interface screen, with instructions printed at the top of the screen. For each prediction trial, a depiction of a light in the middle of the screen lit up with the color red or green randomly on each trial according to the specified probabilities. The layout of the interface screens (with sample feedback) is depicted in the [Appendix](#), for both the single-trial (Fig. 7) and sequence (Fig. 8) conditions. The right–left positions of the two response buttons (red or green) and the probabilities of each color were randomly assigned for each individual participant. Participants were informed only that which light occurred was determined randomly, and they were given the relevant probabilities of each alternative—for example, “The light below will randomly light up with the color Green or Red. About 50 % of the time it will show Red, and 50 % of the time it will show Green.”

In the single-trial prediction conditions (single-trial_bias, single-trial_no-bias), a participant experienced 100 individual trials, making a prediction and receiving feedback on each trial. A bonus of \$0.01 was given for each correct prediction; thus, each participant in the no-bias (50 %–50 %) condition could expect to earn approximately $(.5)(.01)(100) = \$0.50$ in bonus pay, regardless of strategy (see Table 1). In the bias (75 %–25 %) single-trial condition, the expected payoff was $[(.75)(.75) + (.25)(.25)](.01)(100) = \0.625 for a probability-matching strategy (i.e., on each trial, the participant guesses the more likely option with probability .75 and the less likely

Table 1 Study 1: Expected payoffs for consistent use of matching versus maximizing, across conditions

Condition		Strategy	
		Matching	Maximizing
No bias (50 %–50 %)	Single-trial	\$0.50	\$0.50
	Sequence	\$0.156	\$0.156
Bias (75 %–25 %)	Single-trial	\$0.625	\$0.75
	Sequence	\$0.381	\$0.791

option with probability .25), and $[(1)(.75)](.01)(100) = \$0.75$ under maximizing (always predicting the more likely outcome).

In the four-trial sequence condition, 100 trials were also presented to participants, as 25 four-trial sets. Participants were informed that they would earn the bonus only if they correctly predicted all four trials in the sequence. In this condition, a set of four boxes near the bottom of the screen represented the four trials, and the participant clicked on the red or green button for each trial, making predictions for the four trials sequentially. Feedback for individual trials within a sequence was also given sequentially, with a lag of about 0.3 s preceding the appearance of each successive trial's feedback. A bonus of \$0.10 was given for correct prediction of each sequence; thus, participants in the no-bias (50 %–50 %) condition could expect to earn approximately $(.5^4)(.10)(25) = \$0.156$ in bonus pay, regardless of strategy. The expected bonus pay in the bias (75 %–25 %) condition was $(.625^4)(.10)(25) = (.1526)(.10)(25) = \0.381 for probability matching and $(.75^4)(.10)(25) = (.3164)(.10)(25) = \0.791 for maximizing.

It is relevant to note that these rewards approximately equated the expected payoffs for consistent maximizing in the bias (75 %–25 %) condition: \$0.75 for a participant in the single-trial condition versus \$0.79 for a participant in the sequence condition. However, the *gain* in expected total payoff realized by consistent choice of a maximizing versus a matching strategy was $\$0.75 - \$0.625 = \$0.125$ for the single-trial condition and $\$0.791 - \$0.381 = \$0.41$ for the sequence condition. Thus, purely monetary considerations would provide a stronger incentive for choosing to maximize in the sequence condition, opposite to the hypothesized effect of the perfect-prediction goal assumed to be induced by the reward contingencies of the sequence condition.

Results

We coded the apparent strategy used by a participant in each set of four trials, regardless of condition. In the bias (75 %) condition, we scored the responses as *maximizing* if all four predictions were the more likely option, and as *probability matching* if exactly three out of the four predictions were the more likely option. In the no-bias (50 %–50 %) condition, we scored the responses as probability matching if the participant generated a mixed sequence with exactly two predictions of each outcome. Graphs of the proportions of matching across sets for the no-bias and bias conditions are given in Figs. 1 and 2.

Figure 1 shows that in the no-bias (50 %) condition, probability-matching behavior was more prevalent in the sequence condition (where it was observed in 48.8 % of four-

Fig. 1 Proportions of participants using a probability-matching strategy, for the no-bias 50 %–50 % conditions (single-trial and sequence) of Study 1, and for the grouped single-trial condition of Study 1B.

trial sets) than in the single-trial condition (33.7 %). This difference was confirmed in a repeated measures analysis of variance (ANOVA) comparing the two conditions. For this analysis, the mean proportion of matching behavior was computed across five “blocks,” each of which consisted of five four-trial sets (i.e., 20 trials). The main effect of goal type (sequence or single-trial) was significant, $F(1, 173) = 28.01$, $p < .001$, $\eta_p^2 = .140$. This result confirmed our hypothesis that when perfect prediction of a sequence of outcomes is the explicitly rewarded goal, people are more likely to match their predictions in proportion to the probability of the random events. The effects of block and the Block $\times$ Goal Type interaction were not significant, $F(4, 688) = 0.631$, $p = .635$ (with Huynh–Feldt correction), $\eta_p^2 = .004$, and $F(4, 688) = 1.461$, $p = .214$ (with Huynh–Feldt correction), $\eta_p^2 = .008$, respectively.

Figure 2 shows that in the bias (75 %–25 %) condition, the main prediction was again confirmed: Respondents in the Study 1 sequence condition exhibited more probability matching (in 37.8 % of the four-trial sets) than did participants in the single-trial condition (26.9 %). This effect was confirmed in a repeated measures ANOVA, in which the main effect of goal type was significant, $F(1, 173) = 11.22$, $p <$

Fig. 2 Proportions of probability-matching responses across sets in the bias (75 %–25 %) conditions of Studies 1 and 1B.

.001, $\eta_p^2 = .061$. The main effect of block was also significant, $F(4, 692) = 11.66$, $p < .001$, $\eta_p^2 = .063$ (with Huynh–Feldt correction), as was the Block $\times$ Goal Type interaction, $F(4, 692) = 2.789$, $p = .029$ (with Huynh–Feldt correction), $\eta_p^2 = .016$.

The significant effect of block seems to have been due to a trend for probability matching to become less prevalent over the 25 sets of trials. This linear trend was significant, $F(1, 173) = 28.80$, $p < .001$, $\eta_p^2 = .143$, although the slopes did differ between the sequence and single-trial conditions, $F(1, 173) = 6.81$, $p = .010$, $\eta_p^2 = .038$. The decline in probability-matching behavior across trials in the bias condition was accompanied by an increase in probability maximizing, as is shown in Fig. 3. Across blocks, we observed a highly significant linear trend, $F(1, 173) = 104.72$, $p < .001$, $\eta_p^2 = .377$, although the slopes differed slightly between the sequence and single-trial conditions, $F(1, 173) = 3.25$, $p = .073$, $\eta_p^2 = .018$. On average, respondents in the single-trial condition exhibited marginally more maximizing behavior than did participants in the sequence condition, $F(1, 173) = 3.532$, $p = .062$, $\eta_p^2 = .020$.

We also analyzed the simple probability of selecting the more likely outcome, since this is the outcome variable that has been most often reported in previous studies of probability matching. The decline in matching behavior and the (nonsignificant) increase in maximizing across sets of the bias (75 %–25 %) condition result in a gradually increasing probability of selecting the more likely outcome. In the no-bias (50 %–50 %) condition, participants emitted a roughly equiprobable mix of red–green guesses across the 100 trials. In the bias (75 %–25 %) condition, participants started out exhibiting matching behavior, emitting just under 75 % predictions of the more likely outcome. But over the course of the 100 trials, maximizing behavior became more prevalent, culminating in approximately 85 % predictions of the more likely event. This result, that the probability of predicting the more likely outcomes lay above the level predicted by probability matching, has been well documented in the probability-

learning literature (e.g., Baum, 1979; Goldstone & Ashpole, 2004; Myers, 1976) and has been termed “overmatching.”

The observed increase in maximizing behavior with repeated trials corroborates previous results by Newell and Rakow (2007) and Newell, Koehler, James, Rakow, and Ravenzwaaij (2013). Such an increase is especially interesting because this was a description-based task; that is, participants were aware of the outcome probabilities before responding to the first trial (cf. Chen & Corter, 2014). Consistent with Newell et al. (2013), we interpreted this trend as indicating strategy learning in our participants, wherein the experience of outcome feedback led to higher availability, and thus more use, of the maximizing strategy.

This interpretation is supported by the graphs presented in Fig. 4. These graphs show the distribution of inferred strategies (maximizing, minimizing, etc.), as indicated by the proportions of choices of the more likely option for each set of four trials, in the first five sets and the last five sets of trials in the bias (75 %) condition. As we had also assumed above, an observed proportion of .75 choices of the more likely option for a sequence or set of four trials was taken as indicating probability matching, and a proportion of 1 defined the maximizing response, whereas a .5 proportion indicated a “naïve” 50 %–50 % prediction strategy. In the single-trial conditions (top panels), maximizing was the most common response in both the initial 20 trials and the final 20 trials. In contrast, in the sequence condition (lower panels), matching was the modal response in the first five sequences, but maximizing was most common by the final five sequences.

Discussion

Study 1 demonstrated that participants exhibit more probability-matching behavior when the task is presented as the prediction of sequences rather than of individual trials and when perfect prediction of a sequence is required to earn a performance bonus. The bonus for perfect prediction of a four-trial sequence presumably affects the goals that participants adopt in the task. It should be emphasized that these results are in the opposite direction from what might be expected from purely monetary considerations, since the sequence condition here more strongly incentivized a switch from matching to maximizing (as we noted in the Method section).

The results of Study 1 seemed to confirm the hypothesis that explicit or implicit adoption of a perfect-prediction goal for sequences is one cause of probability-matching behavior. However, the single-trial and sequence conditions of this study differed in two important respects: Not only did the sequence condition manipulate goals by rewarding perfect prediction of four-trial sequences, it also grouped the predictions into four-trial sequences, making these strings of observed and predicted outcomes more salient. As we

Fig. 3 Proportions of maximizing responses across sets in the bias (75 %–25 %) conditions of Studies 1 and 1B.

Fig. 4 Distribution of inferred strategies, as indicated by the proportions of choices of the more likely outcome in each four-trial set, over the first five sets of trials (left panels) and the last five sets of trials (right panels).

The top panels show the single-trial bias (75 %) condition; the bottom panels show the sequence bias (75 %) condition.

mentioned, James and Koehler (2011) have demonstrated that an increased emphasis on individual trials rather than sequences of trials can suppress probability matching and encourage maximizing. Thus, it may be that the display of trials and outcomes in four-trial sets in the sequence condition was what drove the increased level of probability matching, rather than the induction of a perfect-prediction goal via the reward structure.

To investigate this possibility, we designed both a follow-up control condition for Study 1 (described below) and a

second study (Study 2) in which we manipulated goal type by varying the reward scheme, holding constant all other aspects of the task.

Study 1B

As we discussed above, the design of Study 1 confounded reward contingency (a large reward for perfect prediction across a sequence of four trials vs. a small reward for individual-trial predictions) with mode of stimulus

presentation (simultaneous presentation of four trials per screen vs. the presentation of individual trials). Accordingly, a follow-up control condition was designed and run, wherein single outcome–feedback trials were presented using the same screen format mode as in the Study 1 sequence condition. In other words, participants were rewarded for correct single-trial predictions, not perfect prediction of sequences, and feedback and payoff information was given after the prediction of each *single* trial. However, here four single trials were shown side by side, as in the sequence condition of Study 1. We will refer to this procedure as the “grouped single-trial” condition.

If we are correct in hypothesizing that the reward structure is what induced the perfect-prediction goals that led to greater matching in the Study 1 sequence condition, then this follow-up control condition should result in performance resembling that in the Study 1 single-trial condition. On the other hand, if the confounding factor of grouped versus single-trial presentation was driving the Study 1 results, then in this follow-up study, participants’ prediction behavior should more closely resemble that for the sequence condition in Study 1.

Method

Participants The participants were again recruited from Amazon’s Mechanical Turk, with the same sample restrictions (i.e., US workers and approval rate of at least 95 %). In this sample ($N = 200$), 58 % of the participants were male, with an average age of 31.9 years, ranging from 19 to 72. Roughly 47 % had taken one to three college-level mathematics or statistics courses, and 5 % had more than three such courses.

Procedure As in Study 1, the interface program was developed in the HTML language with embedded JavaScript. Upon starting the task, participants were randomly assigned to one of the two conditions (bias or no bias). They were then shown the same instructions used for the single-trial condition of Study 1. These instructions stipulated that the color of the light on each trial, red or green, was to be determined randomly, and they specified the relevant probabilities for the two colors. The layout of the interface screen (with sample feedback) is depicted in the Appendix, Fig. 9. The right–left positions of the two response buttons (red or green) and the probabilities of each color were counterbalanced. Participants experienced 100 individual trials (presented as 25 four-trial sets), making a prediction and receiving feedback on each trial. A bonus of \$0.01 was given for each correct prediction; thus, the expected payoffs under matching and maximizing strategies were identical to those in the Study 1 single-trial condition.

Results

The proportions of participants exhibiting probability matching across sets of trials in Study 1B are shown in Figs. 1 and 2, along with the Study 1 results. In the no-bias (50 %–50 %) condition (Fig. 1), the proportion of matching behavior (36.8 %) was similar to that in the Study 1 single-trial condition (33.7 %), and was lower than that in the Study 1 sequence condition (48.8 %). Differences in performance among the three experimental conditions of Studies 1 and 1B were compared in a single repeated measures ANOVA, in which the mean proportions of matching behavior were computed for each of five “blocks,” each consisting of five four-trial sets (i.e., 20 trials). In this analysis, the main effect of condition was significant, $F(2, 271) = 10.74$, $p < .001$, $\eta_p^2 = .073$. Post-hoc tests (with Bonferroni correction) were used to assess the pairwise differences among the three conditions. For the no-bias condition, the proportion of matching for the grouped single-trial condition (Study 1B) was significantly lower than in the Study 1 sequence condition, $F(1, 272) = 12.841$, $p < .001$, $\eta_p^2 = .047$, and was not significantly different from that in the Study 1 single-trial condition, $F(1, 272) = 1.082$, $p = 1.00$, $\eta_p^2 = .003$. This pattern of performance strongly favors the perfect-goal interpretation of the Study 1 results.

For the bias (75 %–25 %) condition of Study 1B (see Fig. 2), the three conditions differed with respect to their overall levels of matching, $F(2, 272) = 14.307$, $p < .001$, $\eta_p^2 = .095$. Pairwise comparisons (using a Bonferroni correction) showed that respondents in the grouped single-trial condition showed less probability matching overall than did respondents in the sequence condition of Study 1 (20.0 % vs. 37.8 %), $F(1, 272) = 28.368$, $p < .001$, $\eta_p^2 = .104$, but did not differ significantly from the Study 1 single-trial participants in their rate of matching (at 26.9 %), $F(1, 272) = 4.189$, $p = .126$, $\eta_p^2 = .015$. Again, these findings support the perfect-goal interpretation of the Study 1 results and do not favor an explanation based on confounding with grouped presentation of trials.

In the original Study 1 bias (75 %–25 %) conditions, the proportion of maximizing behavior was higher in the single-trial condition than in the sequence condition. Here, too, maximizing behavior was more prevalent for the grouped single-trial condition than in the Study 1 sequence condition (Fig. 3), $F(1, 272) = 28.72$, $p < .001$, $\eta_p^2 = .106$. In fact, the level of maximizing behavior for the grouped single-trial condition exceeded that for the Study 1 single-trial condition, $F(1, 272) = 12.49$, $p < .001$, $\eta_p^2 = .046$. Interestingly, this advantage was observed for all four-trial sets except the very first, suggesting that feedback-based *learning* of maximizing strategies (cf. Newell et al., 2013) was facilitated by the presentation of grouped trials.

Discussion

For the no-bias (50 %–50 %) condition of Study 1B, the level of probability matching in the grouped single-trial condition essentially mimicked the behavior in the single-trial condition of Study 1, providing strong support for the perfect-goal hypothesis over an alternative account based on grouping of trials. For the bias (75 %–25 %) condition, the grouped single-trial results did not differ from those in the Study 1 single-trial condition in their level of probability matching, again logically eliminating the confounding Group Trials factor as a cause of the *higher* level of probability matching found for the sequence condition of Study 1.

The (unexpected) higher level of maximizing found for the grouped single-trial condition for every set of trials beyond the initial one (see Fig. 3) suggests that learning of the maximization strategy might be facilitated not only by realistic “satisficing” performance goals, but also by simultaneous displays of multiple trial outcomes. This finding has precedence in the literature on the learning of rule-based and relational abstractions (e.g., Christie & Gentner, 2010), suggesting that the learning of such concepts is facilitated by simultaneous display of their exemplars or instances. It also underscores the conclusion of Newell et al. (2013) that the effects of feedback on the learning of maximizing behavior can vary, depending on individual differences and the task parameters.

Study 2

In Study 2, we investigated whether setting an unrealistically high goal can lead to more probability-matching behavior than does setting a goal that is more realistic. The experimental factor manipulated between subjects was Goal Level, which refers to three levels of performance that participants needed to achieve in order to earn bonus pay: “perfect goal” (100 % correct predictions in a set of five trials), “medium goal” (at least 80 % correct), and “low goal” (at least 60 % correct). Our hypothesis was that probability-matching behavior would be more prevalent in the perfect-goal condition than in the low-goal and medium-goal conditions.

Method

Participants The participants ($N = 300$) were recruited using Amazon’s Mechanical Turk with the same sample restrictions as in the previous two studies (i.e., US workers and approval rate of at least 95 %). Those who had participated previously were blocked from participating in Study 2. The mean age of the participants was 30.70 years, with a range from 17 to 68.

The sample was 61 % male and 39 % female. English was the native language for 70.6 % of our participants.

Procedure The interface used for Study 2 was essentially that used in the sequence condition of Study 1. The task was to predict five-trial sequences of a random binary sequence (the color of a light), with outcome feedback after each sequence. On each trial, the probability of one color showing was .8.

Goals were manipulated by specifying different pay-off policies to different groups of participants. The conditions differed only in the performance goal that would lead to a bonus payment. In the “low-goal” condition, participants need to correctly predict at least three out of five trials (60 %) to earn the performance bonus; in the “medium-goal” condition, the rewarded level was at least four out of five (80 %) correct; in the “perfect-goal” condition, five out of five (100 %) correct predictions were required. In all conditions, participants received a \$0.25 bonus (equal to the base pay that they earned) if they met the specified prediction goal. Participants were randomly assigned to one of the three conditions. After the prediction task, all participants were asked to complete a questionnaire asking about their goals and strategies, as a manipulation check.

In all conditions, maximizing had a higher expected payout than matching (see Table 2). The probability of correctly predicting a single trial outcome was $(.8)(.8) + (.2)(.2) = .68$ under matching and $(1)(.8) = .8$ under maximizing. The expected payout for a block of $n = 5$ trials under each strategy was then given by the appropriate upper-tail cumulative binomial distribution for the prediction goal level times the reward (\$0.25) for reaching the goal. Accordingly, for the low-goal condition the difference in the expected payoffs for maximizing versus matching was $\$1.18 - \$1.01 = \$0.17$; for the medium-goal condition it was $\$0.92 - \$0.61 = \$0.31$; and for the (perfect) high-goal condition it was $\$0.41 - \$0.18 = \$0.23$. Thus, the difference in the expected payoffs between maximizing and matching (one possible index of incentive to maximize) was intermediate for the perfect-goal condition, between the values for the low- and medium-goal conditions.

Table 2 Study 2: Expected payoffs for consistent use of matching versus maximizing, across conditions

Goal Condition	Strategy	
	Matching	Maximizing
Low	\$1.01	\$1.18
Medium	\$0.61	\$0.92
Perfect	\$0.18	\$0.41

Results

Figure 5 shows the proportions of participants in each set of trials who showed probability matching, as indicated by a prediction sequence with exactly four out of five predictions of the more likely (80 %) event. It can be seen that participants exhibited more probability matching in the perfect-goal condition (on 48.4 % of sets) than in the medium- and low-goal conditions (37.0 % and 34.4 %, respectively). This difference was confirmed by a test of the corresponding Helmert contrast in a repeated measures ANOVA with the factors Goal Condition and Set (of trials), which was significant, $F(1, 297) = 9.841$, $p = .002$, $\eta_p^2 = .032$. The observed proportions of probability matching did not differ between the medium- and low-goal conditions, $F(1, 297) = 0.309$, $p = .578$, $\eta_p^2 = .001$.

Figure 6 suggests that the higher prevalence of probability matching in the perfect-goal condition coincided with a slightly lower proportion of maximizing behavior for this condition, at least in later sets. However, the main effect of goal condition for maximizing across all blocks was not significant, $F(2, 297) = 0.086$, $p = .918$, $\eta_p^2 = .001$.

Discussion

Study 2 provided confirmation of the idea that the adoption of a perfect-prediction goal tends to elicit more probability-matching behavior, confirming and extending the results of Study 1. Because the goal conditions in Study 2 did not differ in how they emphasized sequence-level versus single-trial predictions, these findings provide more direct support for the hypothesis that adopting the unrealistic goal of perfect prediction of sequences is a cause of probability matching in probabilistic prediction tasks. As in Study 1, the higher level of probability matching observed in the perfect-goal condition cannot be explained by monetary incentives, since the proportional increase in expected payoff that could result from

Fig. 5 Proportions of participants using a *matching* strategy, across sets, in the three goal-level conditions of Study 2.

Fig. 6 Proportions of participants using a *maximizing* strategy, across sets, in the three goal-level conditions of Study 2.

selecting a maximizing strategy was highest for the perfect-goal condition.

General discussion

The present results demonstrate that emphasizing and rewarding perfect prediction of sequences leads to more use of the “irrational” probability-matching strategy, confirming an idea that has been the subject of prior speculation in the probability-matching literature (e.g., Arkes et al., 1986; Estes, 1964; Fantino & Esfandiari, 2002; Weisberg, 1980). Interestingly, although in Fantino and Esfandiari’s study the main effect of goal instruction was not significant, a specific condition in which their goal instruction seemed to have the expected effect was one in which participants were informed in advance of the relevant outcome probabilities, as in the present studies. Thus, Fantino and Esfandiari’s results are not inconsistent with the present findings.

The idea that adopting a perfect-prediction goal might generally lead to suboptimal performance seems somewhat paradoxical, although it calls to mind Voltaire’s dictum that “The perfect is the enemy of the good.” It has been recognized (e.g., by West & Stanovich, 2003) that there are tasks in which a matching strategy is in fact optimal—for example, predicting the sequence of colors of cards drawn *without replacement* from a small deck (consisting of, say, seven red and three black cards). In this case, perfect performance across the entire deck can *only* result from a matching strategy (although even here more single-trial hits can be expected to result from predicting all “red”). Thus, it is interesting to speculate that the use of the matching strategy may be a learned behavior, reinforced by real-world experience in situations involving sampling without replacement (cf. Ayton & Fischer, 2004; Estes, 1964). Differences in the levels of matching observed in problems invoking various real-world scenarios, such as drawing balls from an urn versus rolling a die (Gal & Baron, 1996) or independent events versus “pick a winner” scenarios (Gao & Corter, 2014), lend some credence to this idea.

It might be seen as a limitation that our procedures and tasks were somewhat atypical of those in probability-matching studies, even among those that have made outcome probabilities known to participants from the outset, because most previous probability-matching studies have used a single-trial format. However, we argue that participants in these typical studies do perceive sequences of outcomes and do think beyond the “frame” of single trials. This assumption is clearly shared by those researchers who believe that matching is the result of higher-order cognitive processes, such as a search for patterns. And, of course, thinking of sequences is a necessary prerequisite of implicitly adopting perfect goals for sequence prediction (although the results of Study 1B show that mere grouping of trials is not sufficient to induce perfect-prediction goals).

Our results are consistent with recent findings by James and Koehler (2011). In two experiments, they found that having participants make predictions for ten unique games, all with equivalent outcome probabilities, led to less matching and more maximization than did making predictions for ten repeated trials of a single game. Their interpretation was that the individuating features of the separate games made it less likely that participants would generate and apply sequence-wide expectations in the prediction task. In a third experiment, in a “global-focus” condition, they encouraged the generation of sequence-level expectations; in a contrasting “local-focus” condition, they emphasized individual trials. The local-focus condition led to more maximizing and less matching.

Thus, our Study 1 findings complement James and Koehler’s (2011) results: James and Koehler achieved an increase in maximizing by emphasizing individual trials, and our Study 1 demonstrated a decrease in maximizing/increase in matching by emphasizing sequences (and perfect goals for the sequences). Furthermore, Study 1 extended James and Koehler’s empirical results in several specific ways. First, the participants in Study 1 received actual monetary rewards for successful performance, lending confidence to the inference that they were unaware of or genuinely confused about the optimal strategy, rather than merely indifferent. Second, James and Koehler used only a bias condition (typical for probability-matching studies), whereas in our Study 1 increased probability-matching behavior was observed for sequences in both the no-bias (50 %–50 %) condition and the bias (75 %–25 %) condition.

However, the results of our three studies taken together show that *presenting* trials grouped into short sequences is not what promotes matching; rather, the adoption of a perfect-prediction goal for sequences (here, induced by the reward structure) seems to be most effective for increasing matching behavior. In fact, highlighting sequences while promoting more attainable goals by rewarding only single-trial predictions (Study 1B) did not especially promote matching, but it did boost maximizing in the bias (75 %–25 %) condition.

Thus, the perfect-goal hypothesis, coupled with the speculative suggestion that participants might adopt *implicit* perfect-prediction goals in some tasks, offers a possible explanation for why (and when) an emphasis on sequences can promote matching, as was proposed by James and Koehler (2011).

Importantly, in the no-bias condition of Study 1 the higher level of matching in the sequence condition could not be attributed to differences in monetary incentives (because the expected payoffs were equal for all prediction strategies), nor to differences in the speeds of learning to maximize—for example, via reinforcement learning—because there was no unique “maximizing” strategy. Rather, higher matching seems to reflect relatively stable (and, in certain circumstances, “irrational”) differences in prediction strategy (e.g., Bar-Hillel, 1982) when participants focus on correct (perfect) prediction of sequences rather than on single-trial outcomes. Furthermore, monetary incentives cannot explain the perfect-goal effect for the bias (75 %–25 %) sequence condition of Study 1, because the difference in the expected payoffs between maximizing and matching was highest for the sequence condition, yet that condition led to the lowest level of maximizing and the highest level of matching. Our Study 2 results most clearly confirmed the hypothesis that performance goals can play an important role in linking sequence-level beliefs and expectations to behavior, because only this factor was manipulated across conditions.

The present results support arguments that goals, and variations in goals, play a critical role in human decision making (Krantz & Kunreuther, 2007; Kruglanski et al., 2002; Zhang, Fishbach, & Dhar, 2007). In our studies, we explicitly manipulated rewards in order to promote the adoption of a perfect-prediction goal. But, as was noted above, we believe that often people implicitly adopt perfect-prediction goals in sequential prediction tasks, even in the absence of explicit rewards. This hypothesis should be more directly investigated in future research. In the meantime, we note that the idea seems consistent with a substantial body of previous research on overconfidence in planning and prediction tasks (e.g., Brenner, Koehler, Liberman, & Tversky, 1996; Larrick, Burson, & Soll, 2007; Lichtenstein & Fischhoff, 1977; Merkle & Weber, 2011; Yates, 1990; Yates, Lee, & Shinotsuka, 1996), showing that people often overestimate their ability to predict outcomes in a variety of stochastic and real-world problems. Furthermore, people are willing to formulate goals and make plans on the basis of these overoptimistic predictions (Arkes et al., 1986; Camerer & Lovallo, 1999; Fishbach & Dhar, 2005; Kivetz, Urminsky, & Zheng, 2006; Zhang et al., 2007).

A perfect-prediction goal is certainly optimistic, and it becomes increasingly unrealistic as the length of the prediction string increases. However, we believe that the spontaneous adoption of such unrealistic perfect-prediction goals may be encouraged by a known bias in judgment: namely, the

systematic overestimation of the probabilities of conjunctive events (e.g., Bar-Hillel, 1973). Thus, it could be argued that the phenomenon of probability matching is not an isolated anomaly, but arises from more general and well-known cognitive biases and heuristics (cf. Gal & Baron, 1996; West & Stanovich, 2003).

Finally, our results show that experience with the prediction task (and appropriate feedback) leads to improving performance, as was indicated by group data showing an increase in the use of the optimal maximizing strategy across sets of trials. This finding is consistent with previous empirical investigations (e.g., Myers, 1976; Shanks, Tunney, & McCarthy, 2002; Vulkan, 2000; West & Stanovich, 2003). However, as in the studies by Newell and Rakow (2007) and Newell et al. (2013), in our tasks the probabilities of the two possible outcomes were provided to participants at the outset; therefore, this improvement cannot be ascribed to learning the outcome probabilities. We found shifts in the probability of maximizing even with a small number of trials. We infer that these shifts are due to strategy availability or to some form of strategy learning, as was posited by Newell et al. But we offer the additional suggestion that this shift in strategy may be triggered by a shift in goals, from the relatively unrealistic perfect-prediction goal to a more attainable level of success. Thus, further research may be warranted, not only into how people choose performance goals, but also into how they formulate and apply strategies and improve or abandon them when the strategies prove unsatisfactory for achieving those goals.

Author note J.G. acknowledges the support provided by a Provost's graduate fellowship award from Teachers College during the time that this research was conducted. Portions of this research were described in an unpublished dissertation submitted by J.G. in partial fulfillment of the requirements for a Ph.D. degree, Columbia University, May 2013.

Appendix: User interfaces for Studies 1 and 1B

Fig. 7 Feedback screen for the single-trial condition (Study 1).

Fig. 8 Feedback screen for the sequence condition (Study 1).

Fig. 9 Initial screen for the grouped single-trial condition (Study 1B).

References

- Arkes, H. R., Dawes, R., & Christensen, C. (1986). Factors influencing the use of a decision rule in a probabilistic task. *Organizational Behavior and Human Decision Processes*, 37, 93–110.
- Ayton, P., & Fischer, I. (2004). The hot hand fallacy and the gambler's fallacy: Two faces of subjective randomness? *Memory & Cognition*, 32, 1369–1378. doi:10.3758/BF03206327
- Bar-Hillel, M. (1973). On the subjective probability of compound events. *Organizational Behavior and Human Decision Processes*, 9, 396–406.
- Bar-Hillel, M. (1982). Studies of representativeness. In D. Kahneman, P. Slovic, & A. Tversky (Eds.), *Judgment under uncertainty: Heuristics and biases* (pp. 69–83). Cambridge, UK: Cambridge University Press.
- Bar-Hillel, M., & Wagenaar, W. A. (1991). The perception of randomness. *Advances in Applied Mathematics*, 12, 428–454.
- Baum, W. M. (1979). Matching, undermatching, and overmatching in studies of choice. *Journal of the Experimental Analysis of Behavior*, 32, 269–281.
- Brenner, L. A., Koehler, D. J., Liberman, V., & Tversky, A. (1996). Overconfidence in probability and frequency judgments: A critical examination. *Organizational Behavior and Human Decisions Processes*, 65, 212–219.
- Bush, R., & Mosteller, F. (1955). *Stochastic models for learning*. New York, NY: Wiley.
- Camerer, C. F., & Lovallo, D. (1999). Overconfidence and excess entry: An experimental approach. *American Economic Review*, 89, 306–318.

- Chen, Y. J., & Corter, J. E. (2014). Learning or framing? Effects of outcome feedback on repeated decisions from description. In P. Bello, M. Guarini, M. McShane, & B. Scassellati (Eds.), *Proceedings of the 36th Annual Conference of the Cognitive Science Society* (pp. 397–402). Austin, TX: Cognitive Science Society.
- Christie, S., & Gentner, D. (2010). Where hypotheses come from: Learning new relations by structural alignment. *Journal of Cognition and Development*, 11, 356–373.
- Edwards, W. (1956). Reward probability, amount, and information as determiners of sequential two-alternative decision. *Journal of Experimental Psychology*, 52, 177–188.
- Edwards, W. (1961). Probability learning in 1,000 trials. *Journal of Experimental Psychology*, 4, 385–394.
- Estes, W. K. (1964). Probability learning. In A. W. Melton (Ed.), *Categories of human learning* (pp. 89–128). New York, NY: Academic Press.
- Fantino, E., & Esfandiari, A. (2002). Probability matching: Encouraging optimal responding in humans. *Canadian Journal of Experimental Psychology*, 56, 58–63.
- Fishbach, A., & Dhar, R. (2005). Goals as excuses or guides: The liberating effect of perceived goal progress on choice. *Journal of Consumer Research*, 32, 370–377.
- Fishbach, A., & Ferguson, M. J. (2007). The goal construct in social psychology. In A. W. Kruglanski & E. T. Higgins (Eds.), *Social psychology: Handbook of basic principles* (2nd ed., pp. 490–515). New York, NY: Guilford Press.
- 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 and Reasoning*, 2, 81–98.
- Gao, J., & Corter, J. (2014). Effects of problem schema on successful maximizing in repeated choices. In P. Bello, M. Guarini, M. McShane, & B. Scassellati (Eds.), *Proceedings of the 36th annual conference of the Cognitive Science Society* (pp. 511–516). Austin, TX: Cognitive Science Society.
- Goldstone, R. L., & Ashpole, B. C. (2004). Human foraging behavior in a virtual environment. *Psychonomic Bulletin & Review*, 11, 508–514.
- Hermstein, R. J. (1961). Relative and absolute strength of responses as a function of frequency of reinforcement. *Journal of Experimental Analysis of Behavior*, 4, 267–272. doi:10.1901/jeab.1961.4-267
- Hermstein, R. J., & Loveland, D. H. (1975). Maximizing and matching on concurrent ratio schedules. *Journal of the Experimental Analysis of Behavior*, 24, 107–116.
- James, G., & Koehler, D. (2011). Banking on a bad bet: Probability matching in risky choice is linked to expectation generation. *Psychological Science*, 22, 707–711.
- Kahneman, D., & Tversky, A. (1972). Subjective probability: A judgment of representativeness. *Cognitive Psychology*, 3, 430–454.
- Kivetz, R., Urminsky, O., & Zheng, Y. (2006). The goal-gradient hypothesis resurrected: Purchase acceleration, illusionary goal progress, and customer retention. *Journal of Marketing Research*, 28, 39–58.
- Koehler, D. J., & James, G. (2009). Probability matching in choice under uncertainty: Intuition versus deliberation. *Cognition*, 113, 123–127.
- Koehler, D. J., & James, G. (2010). Probability matching and strategy availability. *Memory & Cognition*, 38, 667–676.
- Kogler, C., & Kuhberger, A. (2007). Dual process theories: A key for understanding the diversification bias? *Journal of Risk and Uncertainty*, 34, 145–154.
- Krantz, D. H., & Kunreuther, H. C. (2007). Goals and plans in decision making. *Judgment and Decision Making*, 2, 137–168.
- Kruglanski, A. W., Shah, J. Y., Fishbach, A., Friedman, R., Chun, W. Y., & Sleeth-Keppler, D. (2002). A theory of goal systems. In M. P. Zanna (Ed.), *Advances in experimental social psychology* (pp. 331–378). San Diego, CA: Academic Press.
- Larrick, R. P., Burson, K. A., & Soll, J. B. (2007). Social comparison and confidence: When thinking you're better than average predicts overconfidence (and when it does not). *Organizational Behavior and Human Performance*, 102, 76–94.
- Lichtenstein, S., & Fischhoff, B. (1977). Do those who know more also know more about how much they know? The calibration of probability judgments. *Organizational Behavior and Human Performance*, 20, 159–183.
- Merkle, & Weber, M. (2011). True overconfidence: The inability of rational information processing to account for apparent overconfidence. *Organizational Behavior and Human Performance*, 116, 262–271.
- Myers, J. (1976). Probability learning and sequence learning. In W. K. Estes (Ed.), *Handbook of learning and cognitive processes: Approaches to human learning and motivation* (pp. 171–205). Hillsdale, NJ: Erlbaum.
- 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. doi:10.3758/s13421-012-0268-3
- Newell, B. R., & Rakow, T. (2007). The role of experience in decisions from description. *Psychonomic Bulletin & Review*, 14, 1133–1139.
- Nowak, M., & Sigmund, K. (1993). A strategy of win–stay, lose–shift that outperforms tit-for-tat in the Prisoner's Dilemma game. *Nature*, 364, 56–58.
- 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, 274–279.
- Shanks, D. R., Tunney, R. J., & McCarthy, J. D. (2002). A reexamination of probability matching and rational choice. *Journal of Behavioral Decision Making*, 15, 233–250.
- Thorndike, E. L. (1898). Animal intelligence: An experimental study of the associative processes in animals. *Psychological Review Monographs Supplement*, 2(8).
- Unturbe, J., & Corominas, J. (2007). Probability matching involves rule-generating ability: A neuropsychological mechanism dealing with probabilities. *Neuropsychology*, 21, 621–630.
- Vitz, P. C. (1964). Preferences for rates of information presented by sequences of tones. *Journal of Experimental Psychology*, 68, 176–183.
- Voss, J. F., Thompson, C. P., & Keegan, J. H. (1959). Acquisition of probabilistic paired associates as a function of S-R1, S-R2 probability. *Journal of Experimental Psychology*, 58, 390–399.
- Vulkan, N. (2000). An economist's perspective on probability matching. *Journal of Economic Surveys*, 14, 101–118. doi:10.1111/1467-6419.00106
- Wagenaar, W. A. (1970). Appreciation of conditional probabilities in binary sequences. *Acta Psychologica*, 34, 348–356.
- Weisberg, R. W. (1980). *Memory, thought, and behavior*. New York, NY: Oxford University Press.
- West, R. F., & Stanovich, K. E. (2003). Is probability matching smart? Associations between probabilistic choices and cognitive ability. *Memory & Cognition*, 31, 243–251. doi:10.3758/BF03194383
- Wolford, G., Miller, M. B., & Gazzaniga, M. (2000). The left hemisphere's role in hypothesis formation. *Journal of Neuroscience*, 20, RC64.
- Wolford, G., Newman, S. E., Miller, M. B., & Wig, G. S. (2004). Searching for patterns in random sequences. *Canadian Journal of Experimental Psychology*, 58, 221–228.
- Yates, J. F. (1990). *Judgment and decision making*. Englewood Cliffs, NJ: Prentice-Hall.
- Yates, J. F., Lee, J., & Shinotsuka, H. (1996). Beliefs about overconfidence, including its cross-national variation. *Organizational Behavior and Human Decision Processes*, 65, 138–147.
- Yellott, J. I., Jr. (1969). Probability learning with noncontingent success. *Journal of Mathematical Psychology*, 6, 541–575.
- Zhang, Y., Fishbach, A., & Dhar, R. (2007). When thinking beats doing: The role of optimistic expectations in goal-based choice. *Journal of Consumer Research*, 34, 567–578.