Choice Strategies in Multiple-Cue Probability Learning

Chris M. White and Derek J. Koehler
University of Waterloo

Choice strategies for selecting among outcomes in multiple-cue probability learning were investigated using a simulated medical diagnosis task. Expected choice probabilities (the proportion of times each outcome was selected given each cue pattern) under alternative choice strategies were constructed from corresponding observed judged probabilities (of each outcome given each cue pattern) and compared with observed choice probabilities. Most of the participants were inferred to have responded by using a deterministic strategy, in which the outcome with the higher judged probability is consistently chosen, rather than a probabilistic strategy, in which an outcome is chosen with a probability equal to its judged probability. Extended practice in the learning environment did not affect choice strategy selection, contrary to reports from previous studies, results of which may instead be attributable to changes with practice in the variability and extremity of the perceived probabilities on which the choices were based.

Keywords: choice strategies, choice probability, judged probability, probability matching, maximizing

Most important choices faced by individuals and organizations must be made under conditions of uncertainty. Generally speaking, the probability of relevant outcomes is not known but instead must be evaluated by the decision maker on the basis of available predictive cues. In making a treatment decision, for example, a physician might have to choose which of several possible diseases is most likely in light of the patient's symptoms. An investment banker, as another example, might have to choose the most promising real estate investment on the basis of key characteristics of properties on the market.

Such choices are based on the perceived probability of possible outcomes. There is a rich body of research regarding how people evaluate the strength of evidence or support for a given outcome in evaluating its probability. Far less research has addressed how choices are made based on these evaluations. Prior research on the strategies that people use to make choices under conditions of uncertainty, reviewed below, has typically used relatively simple experimental environments in which choices must be made without any predictive information other than the prior frequency of each outcome. The current research investigates the choice strategies used in more complex environments in which the likelihood of the outcome of interest must be based on a set of probabilistically predictive cues. Multiple-cue probability learning, in which identification and use of relevant cues is based on previous experience

in the judgment environment, is an important component of many decisions and may involve different choice strategies than those previously identified as involved in very simple predictive tasks. Models of multiple-cue probability learning, furthermore, rely on specification of a choice process, but typically the implemented choice strategy is not tested against alternative specifications.

Two Choice Strategies

Historically, in the earliest investigations of choice strategies, researchers used the basic probability learning paradigm (for reviews of work in this area, see Estes, 1964; Vulkan, 2000). In that paradigm, the participants' task is to predict which of a set of outcomes will occur on each trial. The only predictive information provided is the base rate of each outcome, which must be learned from trial-by-trial outcome feedback. For example, on each trial, either a red light may illuminate with a probability of 0.7 or a blue light with a probability of 0.3; we refer to these as the *objective probabilities*. The *choice probabilities* are the proportions of trials in which a participant predicts that each outcome will occur.

We considered two specific strategies that people could use to make choices: one of a probabilistic type and the other of a deterministic type. When using the probabilistic strategy, one would predict each outcome to occur on the same proportion of trials as one perceives it to have occurred. If perceptions match reality, then the choice probabilities would match the objective probabilities, so this strategy is usually referred to as probability matching. When using the deterministic strategy, one would always predict the outcome to occur that one believes is the most likely to occur. If one's beliefs match reality, then the choice probabilities would be 1 and 0 and would maximize the number of correct choices made, so this strategy is usually referred to as maximizing. We prefer to use the terms *probabilistic choice strategy* and *deterministic choice strategy* rather than *probability matching* and *maximizing* because the former terms describe the type of strategy used, whereas the latter terms describe the data

Chris M. White and Derek J. Koehler, Department of Psychology, University of Waterloo, Waterloo, Ontario, Canada.

Chris M. White is now at HEC (Faculty of Business and Economics), University of Lausanne, Lausanne, Switzerland.

This research was described in Chris M. White's doctoral dissertation. This research was supported by a grant from the Natural Sciences and Engineering Research Council of Canada to Derek J. Koehler. We are grateful to Joseph Baranski, Mike Dixon, Natalia Kotchetova, and Colin MacLeod for helpful comments on Chris M. White's dissertation.

Correspondence concerning this article should be sent to Chris M. White, Bureau 621, HEC Lausanne, Quartier UNIL-Dorigny, Bâtiment Internef, CH-1015 Lausanne, Switzerland. E-mail: christopher.white@unil.ch

expected to be observed when each strategy is used. The expected choice probabilities as a function of objective probabilities under each choice strategy are shown in Figure 1a.

More complex strategies can predict choice probabilities intermediate between those predicted by the probabilistic and deterministic strategies described here (e.g., the logit rule; see Friedman & Massaro, 1998), but we agree with other authors who have argued that “it does not make sense to think of a process intermediate between probabilistic and deterministic” (Estes, 1995, p. 60). Instead, the observation of intermediate choice probabilities can be explained as resulting from the use of a mixture of the two strategies. That is, each strategy could be used on a certain proportion of the trials by a given individual, or a certain proportion of individuals could use each strategy in a consistent manner.

To determine which choice strategy has been used, one must compare the choice probability for an outcome with its perceived probability of occurring. Choice probabilities are expected to equal perceived probabilities under the probabilistic strategy and to be more extreme (i.e., closer to 1 and 0) than perceived probabilities under the deterministic strategy. Traditionally, it is assumed that the perceived probabilities equal the objective probabilities, so the objective probabilities serve as the benchmark for diagnosing choice strategies, as shown in Figure 1a.

Researchers using the basic probability learning paradigm have often found that choice probabilities are approximately equal to the objective probabilities at asymptote (e.g., Estes, 1964), suggesting that a probabilistic choice strategy is used. However, in the basic probability learning paradigm, choice probabilities are sometimes more extreme than objective probabilities at asymptote, particularly when participants receive financial incentives for making accurate responses (e.g., Siegel & Goldstein, 1959; Shanks, Tunney, & McCarthy, 2002) or are given extensive training (e.g., Myers, Fort, Katz, & Suydam, 1963; Shanks et al., 2002). This has also been found in research using the multiple-cue probability learning paradigm (Estes, 1995; Wallsten & Gu, 2003), and again, this is particularly the case when participants receive financial incentives for making accurate responses (Friedman & Massaro, 1998) or are given extensive training (Goodie & Fantino, 1999). These findings suggest that a deterministic choice strategy is at least sometimes used and is encouraged by certain contextual factors.

Comparison Benchmark

The assumption that the objective probabilities are a reasonable approximation of the perceived probabilities has occasionally been tested by asking people to judge the probability of each outcome or

Figure 1. a: Expected choice probabilities assuming that choices are based on the objective probabilities and a deterministic or probabilistic choice strategy. b: Expected choice and judged probabilities assuming that responses are based on reliable but inaccurate perceived probabilities and that choices are based on a deterministic or probabilistic choice strategy. c: Expected choice and judged probabilities assuming that responses are based on unreliable and inaccurate perceived probabilities and that choices are based on a deterministic or probabilistic choice strategy, assuming two different levels of unreliability.

to estimate the frequency with which each outcome has occurred. When this has been done in the basic probability learning paradigm, the judged probabilities and frequency estimates have approximately matched the objective probabilities at asymptote (these judgments have been elicited intermittently throughout learning, Bauer, 1972; Beach, Rose, Sayeki, Wise, & Carter, 1970; or on every trial, Neimark & Shuford, 1959). The assumption that the objective probabilities are a reasonable approximation of the perceived probabilities is therefore well supported in the basic probability learning paradigm.

In the multiple-cue probability learning paradigm, however, when participants judge the probability of a certain outcome occurring given each combination of cues, the judged probabilities display systematic deviations from the objective probabilities even at asymptote. Among other differences, the information conveyed by the cues tends to be overweighted in contrast to that conveyed by the overall prevalence, or base rate, of each outcome (e.g., Gluck & Bower, 1988; Koehler, White, & Grondin, 2003), and the information conveyed by the presence of some cues tends to be overweighted in contrast to that conveyed by the absence of other cues (Koehler, 2000; Koehler et al., 2003; White & Koehler, 2004). In general, the assumption that the objective probabilities are a good approximation of the perceived probabilities has very little support in the multiple-cue probability learning paradigm, which has led to a large number of descriptive models being developed to account for how people arrive at their perceived probabilities in this type of task (e.g., Gluck & Bower, 1988; Koehler et al., 2003; Kruschke, 1992; Medin & Schaffer, 1978; Reed, 1972). Deviations of perceived probabilities from objective probabilities have also been demonstrated in many other research paradigms (e.g., Kahneman & Tversky, 1973, 1982; Novemsky & Kronzon, 1999; Phillips & Edwards, 1966).

In Figure 1b, the choice probabilities are compared directly with judged probabilities as an alternative method of diagnosing choice strategies that does not require assuming, possibly incorrectly, that the perceived probabilities on which choices are based match the objective probabilities. For illustration, we have assumed that the prior probability or base rate of the target outcome is underweighted when judging the probabilities. Judgments and choices are only plotted for the high base-rate outcome, so the judged probabilities are predicted to fall below the objective probabilities.¹ The exact nature of the deviation between the objective probabilities and judged probabilities, however, is not the focus of this research; it is only important that there is some systematic deviation. Figure 1b shows how the choice strategy can be inferred by examining the relation between choice probabilities and judged probabilities, with choice probabilities expected to be more extreme than judged probabilities under the deterministic strategy but not under the probabilistic strategy.

Variability

A further assumption underlying Figures 1a and 1b is that the perceived probabilities are stable from one trial to the next. But there is variability in all cognitive processes, including probability assessments (as discussed by, e.g., Brenner, 2003; Dougherty, 2001; Erev, Wallsten, & Budescu, 1994; Ferrell & McGoey, 1980; Juslin, Olsson, & Björkman, 1997; Soll, 1996). Wallsten and Budescu (1983) suggested that “subjective certainty is not pre-

cisely determined internally, but rather itself has some variability, vagueness, or fuzziness” (p. 167). In the basic probability learning task, variability in perceived likelihood is likely to be minimal because it is based on the same evidence on every trial. However, in the multiple-cue probability learning task, the evidence on which judgments are based differs from one trial to the next. As a result, the perceived probability must be computed (or recalled) anew, possibly resulting in variability in perceived probabilities across trials even when the evidence is held constant. In addition to its effect on judged probabilities, such unreliability would also be expected to affect choice probabilities, causing them to deviate from the expected values of 0 and 1 even under a completely deterministic choice strategy (as discussed by, e.g., Nosofsky & Palmeri, 1997).

Participants in a multiple-cue probability learning task typically encounter the same cue pattern (i.e., configuration of cue values) on multiple trials. We refer to unreliability in probability assessments, holding constant the cue pattern, as *trial-by-trial variability*. Because of this unreliability, even when a deterministic strategy is used, an individual does not always predict the same outcome when presented with the same cue pattern. When the deterministic strategy is used, instead of the choice probabilities always equaling 1 and 0, they become less extreme as the amount of trial-by-trial variability increases (i.e., they regress toward 0.5, assuming that there are only two outcomes). For example, given a particular cue pattern, one may judge that the probability of a given outcome is 0.7 on one occasion, 0.4 on another occasion, 0.9 on another, and 0.6 on yet another. On the basis of these perceived probabilities, the deterministic strategy would yield a choice probability of 0.75 because the perceived probability was greater than 0.5 on 3 of 4 occasions. The probabilistic strategy would on average yield a choice probability of 0.7, equal to the mean of the perceived likelihoods. Figure 1c shows the expected choice probabilities under each choice strategy assuming that the perceived probabilities not only deviate from the objective probabilities but are also unreliable.

Note that, even when the judgments are unreliable, choice probabilities are expected to be more extreme than judged probabilities under the deterministic choice strategy but not under the probabilistic strategy, in which case the choice and judged probabilities should be equally extreme. Therefore, one approach to diagnosing choice strategies is to compare the extremity of choice and judged probabilities. If the choice probabilities are more extreme than the judged probabilities, this suggests that the deterministic strategy is used at least some of the time by some participants. We complemented the measure of extremity with a measure of trial-by-trial variability, which influences the expected difference in extremity between choice and judged probabilities under the deterministic strategy. For the aggregate data, the extremity and variability measures are also useful in identifying systematic changes in choice strategies across experimental conditions or over trials.

The aggregate measures, however, can only tell us whether at least some people used the deterministic strategy some of the time. To diagnose which choice strategy better captures an individual's

¹ The lines would be reflected about the line $y = x$ if the low base-rate outcome were plotted.

data, we used each participant's judged probabilities to generate expected choice probabilities, assuming they exclusively used either the deterministic strategy or the probabilistic strategy. Then, for each participant (and also for the mean data), we determined whether their observed choice probabilities were better fit by those expected under the deterministic or the probabilistic choice strategy.

Experiments 1 and 2

For purposes of generalizability, in Experiment 1, one type of choice (diagnoses) was compared with two types of judgments (probabilities and frequencies) and in Experiment 2, two types of choices (diagnoses and yes–no choices) were compared with one type of judgment (probabilities). Aside from this distinction, Experiments 1 and 2 were similar in their design and results, so we report them together.

Participants completed a multiple-cue probability learning task in which they used a set of four symptoms, each of which was known to be present or absent, to determine which of two possible flu strains a patient had. Diagnoses (i.e., deciding which flu strain the current patient has) were used as one type of choice because they represent the most basic form of choice, known generically as an n -alternative forced choice task. Yes–no choices (i.e., deciding whether the current patient has a designated flu strain) were used as the second type of choice because they more closely resemble the judgment task in that a focal (target) outcome is designated on each trial for the response to be based on. Probability judgments (i.e., judging the probability that the current patient has a designated flu strain) were used as one type of judgment because they are arguably the most direct way to elicit perceived probabilities. Frequency judgments (i.e., estimating the frequency of a designated flu strain among a certain number of patients who have the current symptom pattern) were used as the second type of judgment in light of research suggesting that they may be based on different considerations than are probability judgments (e.g., Gigerenzer, Hoffrage, & Kleinbölting, 1991; Griffin & Tversky, 1992).

Method

Participants. Volunteer participants were recruited from introductory psychology courses at the University of Waterloo, and they participated for course credit. There were 60 participants in Experiment 1 and 91 in Experiment 2. Participants were randomly assigned to one of the six response orderings.

Apparatus. Each participant completed the experiment on a personal computer in an individual room. All stimuli and instructions were displayed on the monitor, and all responses were given on the keyboard.

Procedure. On each trial, a hypothetical patient was presented. The participant was told that the patient had a fever (a symptom that was always present for every patient) plus a certain combination of the other four symptoms. The background color was black and present symptoms were written in white, capital letters (e.g., "COUGH"), whereas absent symptoms were written in green, lowercase letters with the word "no" in front of them (e.g., no cough). Feedback regarding which flu strain the patient had followed the response. Participants did not receive any external incentives to make accurate responses.

Participants in each experiment completed four blocks of 100 trials each. To ensure that participants had some knowledge of the cue–outcome relations prior to making the responses of interest, we elicited no overt response in the first block, which we refer to as the observation block. On each trial in the observation block, participants viewed the symptoms of a hypothetical patient and pressed the *Enter* key when ready to view the flu strain that the patient had. Participants were told that they would need to make responses in subsequent blocks based on what they had learned during the observation block. Reber and Millward (1968) showed that learning and subsequent responses were not affected by whether participants gave an overt response before viewing each outcome. The three experimental blocks, one for each response mode used in the experiment, followed the observation block in a counterbalanced order, yielding six response orderings.

On diagnosis trials, an example of which is shown in Figure 2, participants were asked, "Which of the two flu strains do you think

Figure 2. Example of a diagnosis trial.

this patient has?" and they responded by pressing the *I* or *2* key on the keyboard to signify Flu Strain #1 or #2. Instructions given prior to this block of trials included the following wording: "... you are asked to guess which flu strain the patient in question has." On yes-no trials, participants were asked, "Do you think this patient has ...?" and one of the flu strains was highlighted. They responded by pressing the *Y* or *N* key on the keyboard to signify "yes" or "no." Instructions given prior to this block of trials included the following wording: "... you are asked to guess whether the patient in question has a designated flu strain."

On probability judgment trials, participants were asked, "What is the probability that this patient has ...?" and one of the two flu strains was highlighted. They responded by using the arrow keys to move a box to highlight a probability between 0% (highlighted at the beginning of the trial) and 100% in increments of 10%, and pressing the *Enter* key once the intended box was highlighted. Instructions given prior to this block of trials included the wording: "A judgment of 100% indicates complete certainty that the patient has the designated flu strain. A judgment of 0% indicates complete certainty that the patient does NOT have the designated flu strain. Thus your probability judgments are an indication of how certain you are that the patient has the designated flu strain, with higher ratings indicating greater certainty." On frequency judgment trials, participants were asked, "Out of 10 patients with this set of symptoms, how many would have ...?" and one of the two flu strains was highlighted. Participants responded by using the arrow keys to move a box to highlight a frequency between 0 (highlighted at the beginning of the trial) and 10 in integer values, and pressing the *Enter* key once the intended box was highlighted. Instructions given prior to this block of trials included the following wording: "A judgment of 10 indicates that all 10 patients with the same set of symptoms would have the designated flu strain. A judgment of 0 indicates that none of the 10 patients would have the designated flu strain." Participants were informed that the flu strain designated for judgment on the yes-no, probability judgment, and frequency judgment trials was selected at random on each trial and was not indicative of the correct diagnosis.

Cue structure. The four symptoms (referred to as Cues A–D here) were all probabilistically predictive of the patient's flu strain and were conditionally independent of each other. The top half of Table 1 shows the frequency with which each cue was present or absent in conjunction with each outcome during the 100 trials in each block. The information conveyed by Cues A and B was more diagnostic than that conveyed by Cues C and D. The high and low base-rate outcomes were randomly assigned the labels *Flu Strain #1* and *Flu Strain #2*. One of four symptom names ("COUGH," "EARACHE," "DIZZINESS," and "SNEEZING") was randomly assigned to each cue for each participant, as was the fixed spatial order of the symptoms on the computer screen. The frequency of the cue patterns was determined by the relations between the cues and the outcomes (i.e., the possible cue patterns were sampled representatively). Each cue pattern was presented with the same frequency in each block, and the outcome frequencies were also identical across blocks.

Dependent measures. Choice probabilities and mean judged probabilities were computed for each of the 16 possible symptom patterns encountered by a participant. For the diagnoses and yes-no choices, choice probabilities for each symptom pattern were defined as the proportion of trials involving the symptom

Table 1
Cue Outcome Co-Occurrence Frequencies

Cue	High base-rate outcome		Low base-rate outcome	
	1	0	1	0
<i>Experiments 1 and 2^a</i>				
A	60	15	5	20
B	15	60	20	5
C	51	24	8	17
D	24	51	17	8
<i>Experiment 3^b</i>				
A	51	17	8	24
B	17	51	24	8
C	47	21	10	22
D	21	47	22	10

Note. 1 = cue was present; 0 = cue was absent.

^a High base-rate outcome, $n = 75$; low base-rate outcome, $n = 25$. ^b High base-rate outcome, $n = 68$; low base-rate outcome, $n = 32$.

pattern in which each flu strain was chosen. For the probability and frequency judgments (after rescaling to the unit interval), the mean judgment for each symptom pattern was computed over trials involving that symptom pattern.

The extremity of the choice probabilities or mean judged probabilities was defined as the standard deviation of the choice or mean judged probabilities over the 16 possible symptom patterns (i.e., their variability around the corresponding grand mean). Other measures of extremity (e.g., the mean absolute deviation) yielded similar results. Trial-by-trial variability for the probability and frequency judgments was similarly measured as the standard deviation of the judgments associated with each symptom pattern around the mean judgment for that pattern, aggregated across all possible symptom patterns.

Inferring choice strategies. Expected choice probabilities under the probabilistic and deterministic choice strategies were generated from the judgments separately for each participant.² Under the probabilistic choice strategy, for each symptom pattern the participant's mean judged probability of a flu strain (normalized to sum to 1.0 across the two possible flu strains) was taken as the expected choice probability of the flu strain for that symptom pattern. Under the deterministic choice strategy, the participant is expected to choose a flu strain if its judged probability is greater than 0.5 and should not choose it when its judged probability is less than 0.5. In cases in which the judged probability is exactly 0.5 (which was true for 5.7% of the trials), each flu strain is expected to be chosen with equal probability. Therefore, under the deterministic strategy, for each symptom pattern the expected choice probability of a flu strain was given by the proportion of trials involving the symptom pattern in which the judged probability of the flu strain was greater than 0.5, plus one half of the proportion

² Because no differences were found between the probability and frequency judgments on any of our measures, the data were collapsed across the type of judgment when inferring the choice probabilities.

of trials in which the judgment was exactly 0.5.³ The mean absolute deviation (*MAD*) of each participant's observed choice probabilities from those expected under each choice strategy was computed to determine which strategy better reproduced their actual choices.⁴

Results and Discussion

The data from 3 participants in each of Experiments 1 and 2 were not analyzed because they failed to complete the experiment. The extremity of the choice and judged probabilities was analyzed using a 3 (response mode) $\times$ 6 (response order) mixed-design analysis of variance (ANOVA) for each experiment. There were no main effects or interactions involving response order ($ps > .10$). The mean extremity based on each response mode is shown in Figure 3, in which the error bars represent one standard error of the mean. The response mode affected the extremity of the responses in both experiments: Experiment 1, $F(2, 102) = 51.9, p < .001, MSE = 0.0020$; Experiment 2, $F(2, 164) = 54.4, p < .001, MSE = 0.0016$.⁵ As expected if the deterministic strategy was used at least some of the time, the choice probabilities were generally more extreme than the judged probabilities. In Experiment 1, the diagnoses were more extreme than the probability judgments, $F(1, 102) = 72.4, p < .001$, and the frequency judgments, $F(1, 102) = 80.0, p < .001$. In Experiment 2, the diagnoses were more extreme than the probability judgments, $F(1, 164) = 97.5, p < .001$, as were the yes-no choices, $F(1, 164) = 61.2, p < .001$.

In Experiment 1, the mean trial-by-trial variability of the probability judgments was approximately equal to that of the frequency judgments ($M = 0.208, SEM = 0.010$; $M = 0.201, SEM = 0.009$, respectively), $t(56) = 0.68$. In Experiment 2, the mean variability of the probability judgments was 0.210 ($SEM = 0.009$).

Comparison of the observed choice probabilities to those expected under each choice strategy suggests that most participants used a deterministic choice strategy. In Experiment 1, the observed choice probabilities of 53 of the 57 participants were more closely fit by those expected under the deterministic strategy than by those expected under the probabilistic strategy. In Experiment 2, this was true for 68 of the 88 participants. When the data are aggregated across participants, the aggregate choice probabilities are also better fit by those expected under the deterministic strategy ($MAD = 0.016$) than by those expected under the probabilistic strategy ($MAD = 0.060$).

Figure 4 shows, for each symptom pattern, the observed choice probabilities for the high base-rate outcome (i.e., flu strain) along with those expected under the deterministic and probabilistic strategies, represented in terms of their deviation from the normalized mean judged probability associated with the symptom pattern. Symptom patterns are ordered from the lowest to the highest judged probability. The data are collapsed across Experiments 1 and 2. Choice probabilities are expected to equal mean judged probabilities under the probabilistic strategy, and therefore the expected choice probabilities under this strategy fall along the solid horizontal line representing zero deviation from the mean judged probabilities. Under the deterministic strategy, choice probabilities are generally expected to be lower than judged probabilities for judged probabilities less than 0.5, and to be greater than judged probabilities for judged probabilities greater than 0.5, subject to the variability of the judgments across trials. For the

aggregate data, the observed choice probabilities coincide quite closely with those expected under the deterministic strategy, which is consistent with the fit statistics reported above.

Experiment 3

Researchers who have investigated choice strategies by comparing choice probabilities with objective probabilities have suggested that participants are more likely to adopt a deterministic strategy as they gain more experience with the task (e.g., Goodie & Fantino, 1999; Myers et al., 1963; Shanks et al., 2002). It is possible that such results are attributable to aspects of processing other than the choice strategy changing with experience. Specifically, if the perceived probabilities on which choices are made themselves become more extreme or less variable with experience, choice probabilities become more extreme relative to objective probabilities even in the absence of any change in choice strategy. These two possibilities cannot be distinguished readily unless, as in the present approach, choice probabilities are compared with judged probabilities rather than to objective probabilities. In Experiment 3, we used this approach to investigate possible changes in choice strategy with increasing task experience.

³ The judgments and yes-no choices that participants gave were based on a focal hypothesis that varied randomly between trials. This was not the case for diagnoses. The randomly varying focal hypothesis may have confused participants on some judgment trials and some yes-no choice trials. For example, if one believes that the probability of the high base-rate outcome is 0.8 given the current cue pattern, then on a judgment trial or a yes-no choice trial, one must first determine which is the focal hypothesis before determining the appropriate response. The response mapping is therefore inconsistent from one trial to the next on judgment trials and yes-no choice trials. The response mapping is consistent from one trial to the next on diagnosis trials because there is no focal hypothesis, so one can determine the appropriate response more easily. Tasks involving consistent response mapping are known to be considerably easier than tasks involving inconsistent response mapping (e.g. Schneider & Shiffrin, 1977). Therefore, to compute the expected choice probabilities for the consistently mapped diagnoses from the inconsistently mapped judgments, we removed the data from the judgment trials on which the participants appeared to have been confused as to the response mapping. This was achieved by removing the data from those trials in which the difference between the judgment on the current trial and the mean judged probability for the current cue pattern was more than ± 0.5 . Iteratively removing the data from these trials, thereby ensuring that all remaining judgments were within ± 0.5 of the trimmed mean judgment, resulted in 5.7% of the judgment data being removed in Experiment 1 and 6.0% in Experiment 2. All of the data were used when computing the expected choice probabilities for the yes-no choice trials and when computing the expected choice probabilities in Experiment 3.

⁴ To compute the *MAD*, we weighted the deviations by the number of trials the corresponding cue pattern appeared in. The *MAD* was used because it is not affected by extreme deviations as much as is the more commonly used root mean squared deviation. Such extreme deviations were more likely to exist for the choice probability computed assuming the deterministic strategy than for the choice probabilities computed assuming the probabilistic strategy. This is because the predictions assuming a deterministic strategy are based on inferred binary responses, so these possess less informational content than do the predictions based on a probabilistic strategy, which are based on continuous responses on an 11-point scale.

⁵ The unnormalized, untrimmed data was used for these analyses.

Figure 3. Extremity of the choice or judged probabilities based on each response mode in Experiments 1 and 2. Error bars represent one standard error of the mean.

Participants completed two sessions of 400 trials each on separate days, for a total of 800 trials. Experiments 1 and 2 showed that the differences between choices and judgments do not depend on the type of choice or judgment, so only diagnoses and probability judgments were used in Experiment 3. To ensure that participants gave some of both types of response at each stage of learning, we alternated the response mode every 25 trials, and there were no observation trials.

In Experiments 1 and 2, it was possible that participants were confused as to which was the focal outcome (i.e., the flu strain designated for judgment) on a small proportion of trials (see Footnote 3). To avoid this possibility, in Experiment 3 we held the focal outcome constant from one trial to the next within each subblock of 25 trials but varied it randomly between subblocks. To ensure that participants did not learn the cue structure too quickly, which would make it more difficult to examine effects of learning across blocks, we used a less predictive cue structure in Experiment 3 than we did in Experiments 1 and 2.

In addition to inferring the prevalence with which each choice strategy was used, an instructional manipulation was introduced in the second session of Experiment 3 that was designed to influence the choice strategy adopted by participants. Half of the participants received instructions at the start of and during the second session describing why and how they should use a deterministic strategy. This manipulation had no effect on any of the dependent variables, so the data were collapsed across this variable in the analyses reported here. Experiment 3 also investigated whether participants were aware of which choice strategy they had used to make choices through the addition of three strategy report questions at the end of the experiment. The gender of the participants, their

university grade average, and their *need for cognition* (Cacioppo, Petty, Feinstein, & Jarvis, 1996) were also elicited because gender and intelligence have been suggested by West and Stanovich (2003) to be related to the choice strategy a person adopts in a basic probability learning task. However, none of these individual difference factors or strategy report responses were significantly related to any of our measures, so we do not discuss them further.

Method

Experiment 3 was identical to Experiments 1 and 2 except as described below.

Participants. Volunteer participants ($N = 48$) were recruited from introductory psychology courses at the University of Waterloo, and they participated for course credit.

Procedure. The experiment was divided into two sessions. Most participants completed each of the two sessions on consecutive days; 4 participants had 1 additional intervening day between the sessions. The 400 trials in each session were divided into 16 subblocks of 25 trials. All participants made diagnoses in the odd-numbered subblocks and probability judgments in the even-numbered subblocks.⁶

Cue structure. The bottom half of Table 1 shows the frequency with which the status of each symptom occurred with each

⁶ Which response mode was given first was intended to be counterbalanced across participants. A programming error caused the participants assigned to give probability judgments first to skip the first subblock, thereby causing them to make diagnoses first and to only receive 375 trials per session.

Figure 4. The degree to which the observed and expected choice probabilities deviated from the judged probabilities for each cue pattern in Experiments 1 and 2.

outcome during each repetition of 100 trials within each response mode.

Results and Discussion

The data from 3 participants were not analyzed because they failed to complete the experiment. The data from each session were split into two blocks, each consisting of 100 trials of each response mode. The first and second blocks in the first session are hereafter referred to as Blocks 1 and 2, and the first and second blocks in the second session as Blocks 3 and 4.⁷ At the beginning of the experiment, the participants knew nothing about how each symptom was related to each disease, making it effectively impossible to infer which choice strategy was used in the initial trials. The data from the first 25 trials of each response mode were therefore not included in the analyses, so Block 1 only included 75 trials of each response mode for all participants.

The extremity of the choice and judged probabilities in each block is shown in Figure 5, as is the variability in the judged probabilities. The extremity data were analyzed using a 2 (response mode) $\times$ 4 (block) within-subject ANOVA.⁸ Overall, the choice probabilities were more extreme than the judged probabilities, $F(1, 44) = 135.3$, $p < .001$, $MSE = 0.0096$, suggesting that the deterministic strategy was used at least some of the time. Although extremity increased significantly across the four blocks, $F(2, 106) = 7.28$, $p < .001$, $MSE = 0.0018$, this increase did not differ based on the type of response, that is, the interaction between response mode and block was not significant, $F(2, 102) = 1.37$, $MSE = 0.0014$. This suggests that choice strategies did not change systematically with experience. In addition, the variability

of the judged probabilities decreased significantly across the four blocks, $F(1, 113) = 13.1$, $p < .001$, $MSE = 0.0029$.

Expected choice probabilities under the deterministic and probabilistic strategies were generated, as in Experiments 1 and 2, for each block. The extremity of the expected choice probabilities under each strategy is shown in Figure 5 for comparison with that associated with the observed choice probabilities. The extremity of the choice probabilities is clearly closer to that expected under the deterministic strategy than under the probabilistic strategy. As for the observed choice probabilities, the extremity of the expected choice probabilities under either strategy increased over blocks. This was due both to the increase in extremity and the decrease in variability of the judged probabilities from which the expected choice probabilities were generated. It is important that the increase in the extremity of the observed choice probabilities across blocks was not significantly different from that of the expected choice probabilities under either strategy. That is, the interaction on the extremity variable between source (three levels: observed choice probabilities, expected choice probabilities under the deterministic strategy, and expected choice probabilities under the probabilistic strategy) and block (four levels) was not significant, $F(5, 211) = 0.67$, $MSE = 0.0013$.

⁷ Blocks 2 and 4 actually included only 75 probability judgments for half of the participants because of the programming error described in Footnote 6.

⁸ There were significant violations of sphericity for the block variable, so the Geisser–Greenhouse correction was used, and all degrees of freedom reported are rounded to the nearest integer.

Figure 5. Extremity of the observed and expected choice probabilities and observed judged probabilities, and variability of the observed judged probabilities in each block of Experiment 3.

Individual-level analyses produced results that are in agreement with the group-level analysis. Choice probabilities from the majority of participants were better fit assuming that the deterministic rather than the probabilistic strategy was used. The proportion of participants whose observed choice probabilities were more closely fit by those expected under the deterministic strategy rather than those expected under the probabilistic strategy did not change significantly across the four blocks (32 of the 45 participants in the first block, 37 in the second block, 31 in the third, and 35 in the fourth), $\chi^2(3) = 2.7$, *ns*.

Because the choice probabilities became more extreme with practice, if we had only elicited choice probabilities, we would have concluded that people were more likely to use a deterministic strategy when they had more practice with the task. Because we elicited judgments concurrently with the choices, we were able to show that the increased extremity of the choice probabilities can be accounted for by the fact that the perceived probabilities on which they are based became more extreme and more reliable with practice. Whether one compares choice probabilities with judged probabilities or compares them with objective probabilities determines what conclusions about choice strategy use are drawn from the same data.

Assumption

Before discussing our results in general, we examine a key assumption of our analyses. To generate expected choice probabilities under alternative choice strategies, we assumed that probability and frequency judgments provide a direct reflection of the same perceived likelihoods that the choices are based on. This

strikes us as a reasonable assumption and is consistent with several aspects of our results, such as the close correspondence between observed choice probabilities and the expected choice probabilities generated from the probability judgments assuming a deterministic choice strategy (e.g., Figure 4). Further evidence consistent with the assumption that both choice probabilities and judged probabilities are based on a common evaluation of perceived probability comes from Experiment 3, in which extended practice affected the extremity of the choice probabilities and judged probabilities to a similar degree.

Additional evidence comes from research that we discussed in the introduction, showing that choice probabilities, judged probabilities, and objective probabilities all tended to be approximately equal at asymptote in many basic probability learning experiments (Bauer, 1972; Beach et al., 1970; Neimark & Shuford, 1959). In addition, when researchers have compared probability judgments with probabilities inferred from a series of choices between different bets, the observed and inferred probabilities tended to be approximately equal (e.g., Beach & Phillips, 1967; Beach & Wise, 1969). And in a signal detection task, Egan, Schulman, and Greenberg (1959; see also Green & Swets, 1966) found that estimates of d' based on choices or judgments did not differ significantly.

Many models of probability judgment assume that probability judgments are a direct reflection of perceived probabilities (e.g., Dougherty, Gettys, & Ogden, 1999; Erev et al., 1994; Gigerenzer et al., 1991). On theoretical grounds, Fox and Tversky (1998) argued that evaluations of likelihood tend to precede choice, are useful in predicting choices, and cannot always be derived reliably from choices. As evidence for these claims, they showed that a

Table 2
Base-Rate Utilization

Experiment	Response mode ^a				Inferential statistics			
	Choice		Judgment					
	Diagnosis	Yes–no	Probability	Frequency	<i>dfs</i>	<i>F</i>	<i>p</i>	<i>MSE</i>
1	0.234 (0.028)		0.261 (0.031)	0.249 (0.027)	2, 102	0.32	>.10	0.037
2	0.256 (0.022)	0.238 (0.024)	0.226 (0.021)		2, 164	0.56	>.10	0.034
3	0.210 (0.024)		0.180 (0.025)		1, 44	3.76	>.05	0.021

^a Values are expressed as means (standard errors of the mean).

subadditive bias in probability judgments also influenced choices among prospects.

As far as we are aware, Goodie and Fantino (1999) offered the only evidence inconsistent with the assumption that judgments and choices are based on the same perceived probabilities. They used a single-cue probability learning task in which the two outcomes had different base rates and argued that the information conveyed by the base rates affected choices more than it affected judgments. However, they did not use any objective measures to quantify the degree to which the base-rate information was used nor did they use any statistical tests to investigate the differences between the two response modes. We do both below, using the data from the current experiments.

Other previously established methods of measuring base-rate utilization (e.g., Kruschke & Johansen, 1999; Wallsten & Gu, 2003; Yates, 1982) are based on the quantitative difference between the average probabilities for the high and low base-rate outcomes, but the magnitude of such measures are potentially affected by the choice strategy used. For instance, if choice probabilities and judged probabilities are based on the same perceived probabilities but the choice probabilities are more extreme (as when the deterministic strategy is used), the quantitative difference measure incorrectly gives the impression that the choice probabilities are more responsive than the judged probabilities to the base rate. This problem can be avoided by using an ordinal measure that compares, for each cue pattern, the choice probability or mean judgment for the high base-rate outcome with that for the low base-rate outcome, taking on a value of +1, 0, or -1, where the former is greater than, equal to, or less than the latter, respectively. The mean value of this measure is then taken over all possible cue patterns.⁹ Base-rate utilization as captured by this measure was consistently greater than zero and did not differ between the response modes in any of the three experiments (see Table 2). This is consistent with the assumption that choices and judgments are based on a common evaluation of perceived probability.

General Discussion

We investigated which choice strategies were used in multiple-cue probability learning environments by comparing choice probabilities with judged probabilities in three experiments involving a simulated medical diagnosis task. Choice probabilities were more extreme than judged probabilities. Although most of the choice probabilities were not at the extreme values of 0 or 1, this does not mean that the deterministic (maximizing) strategy could not have

been used. The deterministic strategy only implies choice probabilities to be at the extreme limits of 0 and 1 when there is no variability in perceived probability across repeated presentations of a particular cue pattern. In fact, the comparison to the judged probabilities indicates that the observed choice probabilities of most of the participants in all three experiments more closely resembled those expected under the deterministic rather than under the probabilistic choice strategy. In Experiment 3, choice probabilities became more extreme with extended practice, but this appeared to be a predictable consequence of a decrease in the trial-by-trial variability and an increase in the extremity of the perceived probabilities across trials. It was not necessary to assume that the likelihood of using the deterministic choice strategy changed as participants gained experience in the environment.

Other researchers have demonstrated that variability in perceived probabilities causes judged probabilities to be regressive to the mean (Erev et al., 1994) and therefore less extreme than the corresponding objective probabilities. Although helpful in accounting for other aspects of judgment data, this approach cannot account for our observation of judged probabilities being less extreme than the choice probabilities because the choice probabilities should also be subject to the regression effect. If the probabilistic choice strategy was used to make choices and the choices were based on the same perceived probabilities as the judgments, then the variability would be reflected to the same degree in the choice probabilities as in the judged probabilities, so both would be equally regressive.

Wallsten and Gu (2003) used objective probabilities as the benchmark of comparison for the choice probabilities, but they went beyond what other authors have done with this method by using two free parameters to allow the predicted choice probabilities to deviate from the objective probabilities. One of these parameters had a similar effect on the predicted choice probabil-

⁹ The expected value of this measure is zero when base-rate information is completely ignored; the maximum (minimum) value of +1 (-1) would be achieved by someone who chose the high (low) base-rate outcome most often for every cue pattern, or whose mean judged probability for the high (low) base-rate outcome was above 0.5 for every cue pattern. The objective probability of the high base-rate outcome was higher than that of the low base-rate outcome for 10 of the 16 cue patterns in Experiments 1 and 2 and for 11 of the 16 cue patterns in Experiment 3, so the base-rate utilization of the objective probabilities (i.e., that expected for a Bayesian respondent) was 0.25 in Experiments 1 and 2 and 0.38 in Experiment 3.

ities as did incorporating trial-by-trial variability in the work presented here. That is, it allowed the deterministic strategy to predict choice probabilities that were not at the extreme values of 1 and 0 but instead were somewhat regressive toward the objective probabilities. Their second parameter determined the level of objective probability for which the predicted choice probability would cross 0.5. This had a similar effect as did our use of the judged probabilities to accommodate deviations of perceived probabilities from objective probabilities. The freedom incorporated into Wallsten and Gu's (2003) predictions allowed them to accommodate similar patterns of choice probabilities as were observed here, and they reached similar conclusions, namely that most participants used the deterministic rather than the probabilistic choice strategy. The current research is similar except that we based the expected choice probabilities on the observed judged probabilities, so no assumptions regarding how or why the perceived probabilities deviate from the objective probabilities or vary from trial to trial were necessary.

The current research suggests that much can be learned by comparing choices between outcomes with the judged likelihood of each outcome rather than with the objective likelihood of each outcome. Our results also bring into question the practice of using choice probabilities as direct estimates of perceived probabilities. This is often done when attempting to understand how people perform multiple-cue probability learning tasks, but it is only valid if participants make their choices using the probabilistic choice strategy. The current research suggests that this is not the case for most participants. If models of perceived probability in the multiple-cue probability learning task are to be tested by comparing their predictions with observed choice probabilities, then the models need to first incorporate plausible assumptions regarding choice strategies and consider how the resulting choices are influenced by the variability in the perceived probabilities on which they are based. It may therefore be preferable to test these models by comparing their predictions with judged probabilities rather than with choice probabilities.

References

- Bauer, M. (1972). Relations between prediction- and estimation-responses in cue-probability learning and transfer. *Scandinavian Journal of Psychology*, 13, 198–207.
- Beach, L. R., & Phillips, L. D. (1967). Subjective probabilities inferred from estimates and bets. *Journal of Experimental Psychology*, 75, 354–359.
- Beach, L. R., Rose, R. M., Sayeki, Y., Wise, J. A., & Carter, W. B. (1970). Probability learning: Response proportions and verbal estimates. *Journal of Experimental Psychology*, 86, 165–170.
- Beach, L. R., & Wise, J. A. (1969). Subjective probability revision and subsequent decisions. *Journal of Experimental Psychology*, 81, 561–565.
- Brenner, L. (2003). A random support model of the calibration of subjective probabilities. *Organizational Behavior and Human Decision Processes*, 90, 87–110.
- Cacioppo, J. T., Petty, R. E., Feinstein, J. A., & Jarvis, W. B. G. (1996). Dispositional differences in cognitive motivation: The life and times of individuals varying in need for cognition. *Psychological Bulletin*, 119, 197–253.
- Dougherty, M. R. P. (2001). Integration of the ecological and error models of overconfidence using a multiple-trace memory model. *Journal of Experimental Psychology: General*, 130, 579–599.
- Dougherty, M. R. P., Gettys, C. F., & Ogden, E. E. (1999). Minerva-DM: A memory processes model for judgments of likelihood. *Psychological Review*, 106, 180–209.
- Egan, J. P., Schulman, A. I., & Greenberg, G. Z. (1959). Operating characteristics determined by binary decisions and by ratings. *Journal of the Acoustical Society of America*, 31, 768–773.
- Erev, I., Wallsten, T. S., & Budescu, D. V. (1994). Simultaneous over- and underconfidence: The role of error in judgment processes. *Psychological Review*, 101, 519–527.
- Estes, W. K. (1964). Probability learning. In A. W. Melton (Ed.), *Categories of human learning* (pp. 89–128). New York: Academic Press.
- Estes, W. K. (1995). Response processes in cognitive models. In R. F. Lorch, Jr. & E. J. O'Brien (Eds.), *Sources of coherence in reading* (pp. 51–71). Hillsdale, NJ: Erlbaum.
- Ferrell, W. R., & McGoe, P. J. (1980). A model of calibration for subjective probabilities. *Organizational Behavior and Human Performance*, 26, 32–53.
- Fox, C. R., & Tversky, A. (1998). A belief-based account of decision under uncertainty. *Management Science*, 44, 879–895.
- Friedman, D., & Massaro, D. W. (1998). Understanding variability in binary and continuous choice. *Psychonomic Bulletin and Review*, 5, 370–389.
- Gigerenzer, G., Hoffrage, U., & Kleinbölting, H. (1991). Probabilistic mental models: A Brunswikian theory of confidence. *Psychological Review*, 98, 506–528.
- Gluck, M. A., & Bower, G. H. (1988). From conditioning to category learning: An adaptive network model. *Journal of Experimental Psychology: General*, 117, 227–247.
- Goodie, A. S., & Fantino, E. (1999). What does and does not alleviate base-rate neglect under direct experience. *Journal of Behavioral Decision Making*, 12, 307–335.
- Green, D. M., & Swets, J. A. (1966). *Signal detection theory and psychophysics*. New York: Wiley.
- Griffin, D., & Tversky, A. (1992). The weighing of evidence and the determinants of confidence. *Cognitive Psychology*, 24, 411–435.
- Juslin, P., Olsson, H., & Björkman, M. (1997). Brunswikian and Thurstonian origins of bias in probability assessment: On the interpretation of stochastic components of judgment. *Journal of Behavioral Decision Making*, 10, 189–209.
- Kahneman, D., & Tversky, A. (1973). On the psychology of prediction. *Psychological Review*, 80, 237–251.
- Kahneman, D., & Tversky, A. (1982). On the study of statistical intuitions. *Cognition*, 11, 123–141.
- Koehler, D. J. (2000). Probability judgment in three-category classification learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 26, 28–52.
- Koehler, D. J., White, C. M., & Grondin, R. (2003). An evidential support accumulation model of subjective probability. *Cognitive Psychology*, 46, 152–197.
- Kruschke, J. K. (1992). ALCOVE: An exemplar-based connectionist model of category learning. *Psychological Review*, 99, 22–44.
- Kruschke, J. K., & Johansen, M. K. (1999). A model of probabilistic category learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 25, 1083–1119.
- Medin, D. L., & Schaffer, M. M. (1978). Context theory of classification learning. *Psychological Review*, 85, 207–238.
- Myers, J. L., Fort, J. G., Katz, L., & Suidam, M. M. (1963). Differential monetary gains and losses and event probability in a two-choice situation. *Journal of Experimental Psychology*, 66, 521–522.
- Neimark, E. D., & Shuford, E. H. (1959). Comparison of predictors and estimates in a probability learning situation. *Journal of Experimental Psychology*, 57, 294–298.
- Nosofsky, R. M., & Palmeri, T. J. (1997). An exemplar-based random walk model of speeded classification. *Psychological Review*, 104, 266–300.

- Novemsky, N., & Kronzon, S. (1999). How are base-rates used, when they are used: A comparison of additive and Bayesian models of base-rate use. *Journal of Behavioral Decision Making*, 12, 55–69.
- Phillips, L. D., & Edwards, W. (1966). Conservatism in a simple probability inference task. *Journal of Experimental Psychology*, 72, 346–354.
- Reber, A. S., & Millward, R. B. (1968). Event observation in probability learning. *Journal of Experimental Psychology*, 77, 317–327.
- Reed, S. K. (1972). Pattern recognition and categorization. *Cognitive Psychology*, 3, 383–407.
- Schneider, W., & Shiffrin, R. M. (1977). Controlled and automatic human information processing: I. Detection, search, and attention. *Psychological Review*, 84, 1–66.
- Shanks, D. R., Tunney, R. J., & McCarthy, J. D. (2002). A re-examination of probability matching and rational choice. *Journal of Behavioral Decision Making*, 15, 233–250.
- Siegel, S., & Goldstein, D. A. (1959). Decision-making behavior in a two-choice uncertain outcome situation. *Journal of Experimental Psychology*, 57, 37–42.
- Soll, J. B. (1996). Determinants of overconfidence and miscalibration: The roles of random error and ecological structure. *Organizational Behavior and Human Decision Processes*, 65, 117–137.
- Vulkan, N. (2000). An economist's perspective on probability matching. *Journal of Economic Surveys*, 14, 101–118.
- Wallsten, T. S., & Budescu, D. V. (1983). Encoding subjective probabilities: A psychological and psychometric review. *Management Science*, 29, 151–173.
- Wallsten, T. S., & Gu, H. (2003). Distinguishing choice and subjective probability estimation processes: Implications for theories of judgment and for cross-cultural comparisons. *Organizational Behavior and Human Decision Processes*, 90, 111–123.
- West, R. F., & Stanovich, K. E. (2003). Is probability matching smart? Associations between probabilistic choices and cognitive ability. *Memory & Cognition*, 31, 243–251.
- White, C. M., & Koehler, D. J. (2004). Missing information in multiple-cue probability learning. *Memory and Cognition*, 32, 1007–1018.
- Yates, J. F. (1982). External correspondence: Decompositions of the mean probability score. *Organizational Behavior and Human Performance*, 30, 132–156.

Received October 27, 2006

Revision received March 22, 2007

Accepted April 9, 2007 ■

Members of Underrepresented Groups: Reviewers for Journal Manuscripts Wanted

If you are interested in reviewing manuscripts for APA journals, the APA Publications and Communications Board would like to invite your participation. Manuscript reviewers are vital to the publications process. As a reviewer, you would gain valuable experience in publishing. The P&C Board is particularly interested in encouraging members of underrepresented groups to participate more in this process.

If you are interested in reviewing manuscripts, please write to the address below. Please note the following important points:

- To be selected as a reviewer, you must have published articles in peer-reviewed journals. The experience of publishing provides a reviewer with the basis for preparing a thorough, objective review.
- To be selected, it is critical to be a regular reader of the five to six empirical journals that are most central to the area or journal for which you would like to review. Current knowledge of recently published research provides a reviewer with the knowledge base to evaluate a new submission within the context of existing research.
- To select the appropriate reviewers for each manuscript, the editor needs detailed information. Please include with your letter your vita. In the letter, please identify which APA journal(s) you are interested in, and describe your area of expertise. Be as specific as possible. For example, "social psychology" is not sufficient—you would need to specify "social cognition" or "attitude change" as well.
- Reviewing a manuscript takes time (1–4 hours per manuscript reviewed). If you are selected to review a manuscript, be prepared to invest the necessary time to evaluate the manuscript thoroughly.

Write to Journals Office, American Psychological Association, 750 First Street, NE, Washington, DC 20002-4242.