Journal of Risk and Uncertainty, 4:5-28 (1991)

Preference and Belief: Ambiguity and Competence in Choice under Uncertainty

CHIP HEATH
Stanford University

AMOS TVERSKY
Department of Psychology, Stanford University, Jordan Hall, Bldg. 420, Stanford, CA 94305-2130

Key words: ambiguity, uncertainty, preferences

Abstract

We investigate the relation between judgments of probability and preferences between bets. A series of experiments provides support for the competence hypothesis that people prefer betting on their own judgment over an equiprobable chance event when they consider themselves knowledgeable, but not otherwise. They even pay a significant premium to bet on their judgments. These data cannot be explained by aversion to ambiguity, because judgmental probabilities are more ambiguous than chance events. We interpret the results in terms of the attribution of credit and blame. The possibility of inferring beliefs from preferences is questioned. [1]

The uncertainty we encounter in the world is not readily quantified. We may feel that our favorite football team has a good chance to win the championship match, that the price of gold will probably go up, and that the incumbent mayor is unlikely to be reelected, but we are normally reluctant to assign numerical probabilities to these events. However, to facilitate communication and enhance the analysis of choice, it is often desirable to quantify uncertainty. The most common procedure for quantifying uncertainty involves expressing belief in the language of chance. When we say that the probability of an uncertain event is 30%, for example, we express the belief that this event is as probable as the drawing of a red ball from a box that contains 30 red and 70 green balls. An alternative procedure for measuring subjective probability seeks to infer the degree of belief from preference via expected utility theory. This approach, pioneered by Ramsey (1931) and further developed by Savage (1954) and by Anscombe and Aumann (1963), derives subjective probability from preferences between bets. Specifically, the subjective probability of an uncertain event E is said to be p if the decision maker is

This work was supported by Grant 89-0064 from The Air Force Office of Scientific Research to Stanford University. Funding for experiment 1 was provided by SES 8420240 to Ray Battalio. We have benefited from discussions with Max Bazerman, Daniel Ellsberg, Richard Gonzales, Robin Hogarth, Linda Ginzel, Daniel Kahneman, and Eldar Shafir.

indifferent between the prospect of receiving $\$x$ if E occurs (and nothing otherwise) and the prospect of receiving $\$x$ if a red ball is drawn from a box that contains a proportion p of red balls.

The Ramsey scheme for measuring belief and the theory on which it is based were challenged by Daniel Ellsberg (1961; see also Fellner, 1961) who constructed a compelling demonstration of what has come to be called an ambiguity effect, although the term *vagueness* may be more appropriate. The simplest demonstration of this effect involves two boxes: one contains 50 red balls and 50 green balls, whereas the second contains 100 red and green balls in unknown proportion. You draw a ball blindly from a box and guess its color. If your guess is correct, you win $20; otherwise you get nothing. On which box would you rather bet? Ellsberg argued that people prefer to bet on the 50/50 box rather than on the box with the unknown composition, even though they have no color preferences and so are indifferent between betting on red or on green in either box. This pattern of preferences, which was later confirmed in many experiments, violates the additivity of subjective probability because it implies that the sum of the probabilities of red and of green is higher in the 50/50 box than in the unknown box.

Ellsberg's work has generated a great deal of interest for two reasons. First, it provides an instructive counter example to (subjective) expected utility theory within the context of games of chance. Second, it suggests a general hypothesis that people prefer to bet on clear rather than on vague events, at least for moderate and high probability. For small probability, Ellsberg suggested, people may prefer vagueness to clarity. These observations present a serious problem for expected utility theory and other models of risky choice because, with the notable exception of games of chance, most decisions in the real world depend on uncertain events whose probabilities cannot be precisely assessed. If people's choices depend not only on the degree of uncertainty but also on the precision with which it can be assessed, then the applicability of the standard models of risky choice is severely limited. Indeed, several authors have extended the standard theory by invoking nonadditive measures of belief (e.g., Fishburn, 1988; Schmeidler, 1989) or second-order probability distributions (e.g., Gärdenfors and Sahlin, 1982; Skyrm, 1980) in order to account for the effect of ambiguity. The normative status of these models is a subject of lively debate. Several authors, notably Ellsberg (1963), maintain that aversion to ambiguity can be justified on normative grounds, although Raiffa (1961) has shown that it leads to incoherence.

Ellsberg's example, and most of the subsequent experimental research on the response to ambiguity or vagueness, were confined to chance processes, such as drawing a ball from a box, or problems in which the decision maker is provided with a probability estimate. The potential significance of ambiguity, however, stems from its relevance to the evaluation of evidence in the real world. Is ambiguity aversion limited to games of chance and stated probabilities, or does it also hold for judgmental probabilities? We found no answer to this question in the literature, but there is evidence that casts some doubt on the generality of ambiguity aversion.

For example, Budescu, Weinberg, and Wallsten (1988) compared the cash equivalents given by subjects for gambles whose probabilities were expressed numerically, graphically, or verbally. In the graphical display, probabilities were presented as the shaded area of a circle. In the verbal form, probabilities were described by expressions such as "very likely" or "highly improbable." Because the verbal and the graphical forms are more ambiguous than the numerical form, ambiguity aversion implies a preference for the numerical display. This prediction was not confirmed. Subjects priced the gambles roughly the same in all three displays. In a different experimental paradigm, Cohen and Hansel (1959) and Howell (1971) investigated subjects' choices between compound gambles involving both skill and chance components. For example, in the latter experiment the subject had to hit a target with a dart (where the subjects's hit rate equaled 75%) as well as spin a roulette wheel so that it would land on a marked section composing 40% of the area. Success involves a 75% skill component and 40% chance component with an overall probability of winning of $.75 \times .4 = .3$. Howell varied the skill and chance components of the gambles, holding the overall probability of winning constant. Because the chance level was known to the subject whereas the skill level was not, ambiguity aversion implies that subjects would shift as much uncertainty as possible to the chance component of the gamble. In contrast, 87% of the choices reflect a preference for skill over chance. Cohen and Hansel (1959) obtained essentially the same result.

1. The competence hypothesis

The preceding observations suggest that the aversion to ambiguity observed in a chance setup (involving aleatory uncertainty) does not readily extend to judgmental problems (involving espistemic uncertainty). In this article, we investigate an alternative account of uncertainty preferences, called the competence hypothesis, which applies to both chance and evidential problems. We submit that the willingness to bet on an uncertain event depends not only on the estimated likelihood of that event and the precision of that estimate; it also depends on one's general knowledge or understanding of the relevant context. More specifically, we propose that—holding judged probability constant—people prefer to bet in a context where they consider themselves knowledgeable or competent than in a context where they feel ignorant or uninformed. We assume that our feeling of competence[1] in a given context is determined by what we know relative to what can be known. Thus, it is enhanced by general knowledge, familiarity, and experience, and is diminished, for example, by calling attention to relevant information that is not available to the decision maker, especially if it is available to others.

There are both cognitive and motivational explanations for the competence hypothesis. People may have learned from lifelong experience that they generally do better in situations they understand than in situations where they have less knowledge. This expectation may carry over to situations where the chances of winning are no longer higher in the familiar than in the unfamiliar context. Perhaps the major reason for the competence hypothesis is motivational rather than cognitive. We propose that the consequences of each bet include, besides the monetary payoffs, the credit or blame associated with the outcome. Psychic payoffs of satisfaction or embarrassment can result from self-evaluation or from an evaluation by others. In either case, the credit and the blame associated with an outcome depend, we suggest, on the attributions for success and failure. In the domain of chance, both success and failure are attributed primarily to luck. The situation is different when a person bets on his or her judgment. If the decision maker has limited understanding of the problem at hand, failure will be attributed to ignorance, whereas success is likely to be attributed to chance. In contrast, if the decision maker is an "expert," success is attributable to knowledge, whereas failure can sometimes be attributed to chance.

We do not wish to deny that in situations where experts are supposed to know all the facts, they are probably more embarrassed by failure than are novices. However, in situations that call for an educated guess, experts are sometimes less vulnerable than novices because they can better justify their bets, even if they do not win. In betting on the winner of a football game, for example, people who consider themselves experts can claim credit for a correct prediction and treat an incorrect prediction as an upset. People who do not know much about football, on the other hand, cannot claim much credit for a correct prediction (because they are guessing), and they are exposed to blame for an incorrect prediction (because they are ignorant).

Competence or expertise, therefore, helps people take credit when they succeed and sometimes provides protection against blame when they fail. Ignorance or incompetence, on the other hand, prevents people from taking credit for success and exposes them to blame in case of failure. As a consequence, we propose, the balance of credit to blame is most favorable for bets in one's area of expertise, intermediate for chance events, and least favorable for bets in an area where one has only limited knowledge. This account provides an explanation of the competence hypothesis in terms of the asymmetry of credit and blame induced by knowledge or competence.

The preceding analysis readily applies to Ellsberg's example. People do not like to bet on the unknown box, we suggest, because there is information, namely the proportion of red and green balls in the box, that is knowable in principle but unknown to them. The presence of such data makes people feel less knowledgeable and less competent and reduces the attractiveness of the corresponding bet. A closely related interpretation of Ellsberg's example has been offered by Frisch and Baron (1988). The competence hypothesis is also consistent with the finding of Curley, Yates, and Abrams (1986) that the aversion to ambiguity is enhanced by anticipation that the contents of the unknown box will be shown to others.

Essentially the same analysis applies to the preference for betting on the future rather than on the past. Rothbart and Snyder (1970) asked subjects to roll a die and bet on the outcome either before the die was rolled or after the die was rolled but before the result was revealed. The subjects who predicted the outcome before the die was rolled expressed greater confidence in their guesses than the subjects who predicted the outcome after the die roll ("postdiction"). The former group also bet significantly more money than the latter group. The authors attributed this phenomenon to magical thinking or the illusion of control, namely the belief that one can exercise some control over the outcome before, but not after, the roll of the die. However, the preference to bet on future rather than past events is observed even when the illusion of control does not provide a plausible explanation, as illustrated by the following problem in which subjects were presented with a choice between the two bets:

1. A stock is selected at random from the *Wall Street Journal*. You guess whether it will go up or down tomorrow. If you're right, you win $5.
2. A stock is selected at random from the *Wall Street Journal*. You guess whether it went up or down yesterday. You cannot check the paper. If you're right, you win $5.

Sixty-seven percent of the subjects ($N = 184$) preferred to bet on tomorrow's closing price. (Ten percent of the participants, selected at random, actually played their chosen bet.) Because the past, unlike the future, is knowable in principle, but not to them, subjects prefer the future bet where their relative ignorance is lower. Similarly, Brun and Teigen (1990) observed that subjects preferred to guess the result of a die roll, the sex of a child, or the outcome of a soccer game before the event rather than afterward. Most of the subjects found guessing before the event more "satisfactory if right" and less "uncomfortable if wrong." In prediction, only the future can prove you wrong; in postdiction, you could be wrong right now. The same argument applies to Ellsberg's problem. In the 50/50 box, a guess could turn out to be wrong only after drawing the ball. In the unknown box, on the other hand, the guess may turn out to be mistaken even before the drawing of the ball—if it turns out that the majority of balls in the box are of the opposite color. It is noteworthy that the preference to bet on future rather than on past events cannot be explained in terms of ambiguity because, in these problems, the future is as ambiguous as the past.

Simple chance events, such as drawing a ball from a box with a known composition involve no ambiguity; the chances of winning are known precisely. If betting preferences between equiprobable events are determined by ambiguity, people should prefer to bet on chance over their own vague judgments (at least for moderate and high probability). In contrast, the attributional analysis described above implies that people will prefer betting on their judgment over a matched chance event when they feel knowledgeable and competent, but not otherwise. This prediction is confirmed by the finding that people prefer betting on their skill rather than on chance. It is also consistent with the observation of March and Shapira (1987) that many top managers, who consistently bet on highly uncertain business propositions, resist the analogy between business decisions and games of chance.

We have argued that the present attributional analysis can account for the available evidence on uncertainty preferences, whether or not they involve ambiguity. These include 1) the preference for betting on the known rather than on the unknown box in Ellsberg's problem, 2) the preference to bet on future rather than on past events, and 3) the preference for betting on skill rather than on chance. Furthermore, the competence hypothesis implies a *choice-judgment discrepancy*, namely a preference to bet on A rather than on B even though B is judged to be at least as probable as A. In the following series of experiments, we test the competence hypothesis and investigate the choice-judgment discrepancy. In experiment 1 we offer people the choice between betting on their judged probabilities for general knowledge items or on a matched chance lottery. Experiments 2 and 3 extend the test by studying real-world events and eliciting an independent assessment of knowledge. In experiment 4, we sort subjects according to their area of expertise and compare their willingness to bet on their expert category, a nonexpert category, and chance. Finally, in experiment 5, we test the competence hypothesis in a pricing task that does not involve probability judgment. The relations between belief and preference are discussed in the last section of the article.

1.1. *Experiment 1: Betting on knowledge*

Subjects answered 30 knowledge questions in two different categories, such as history, geography, or sports. Four alternative answers were presented for each question, and the subject first selected a single answer and then rated his or her confidence in that answer on a scale from 25% (pure guessing) to 100% (absolute certainty). Participants were given detailed instructions about the use of the scale and the notion of calibration. Specifically, they were instructed to use the scale so that a confidence rating of 60%, say, would correspond to a hit rate of 60%. They were also told that these ratings would be the basis for a money-making game, and warned that both underconfidence and overconfidence would reduce their earnings.

After answering the questions and assessing confidence, subjects were given an opportunity to choose between betting on their answers or on a lottery in which the probability of winning was equal to their stated confidence. For a confidence rating of 75%, for example, the subject was given the choice between 1) betting that his or her answer was correct, or 2) betting on a 75% lottery, defined by drawing a numbered chip in the range 1–75 from a bag filled with 100 numbered poker chips. For half of the questions, lotteries were directly equated to confidence ratings. For the other half of the questions, subjects chose between the complement of their answer (betting that an answer other than the one they choose is correct) or the complement of their confidence rating. Thus, if a subject chose answer A with confidence of 65%, the subject could choose between betting that one of the remaining answers B, C, or D is correct, or betting on a 100% − 65% = 35% lottery.

Two groups of subjects participated in the experiment. One group (N = 29) included psychology students who received course credit for participation. The second group (N = 26) was recruited from introductory economics classes and performed the experiment for cash earnings. To determine the subjects' payoffs, ten questions were selected at random, and the subjects played out the bets they had chosen. If subjects chose to gamble on their answer, they collected $1.50 if their answer was correct. If subjects chose to bet on the chance lottery, they drew a chip from the bag and collected $1.50 if the number on the chip fell in the proper range. Average earnings for the experiment were around $8.50.

Paid subjects took more time than unpaid subjects in selecting their answers and assessing confidence; they were slightly more accurate. Both groups exhibited overconfidence: the paid subjects answered correctly 47% of the questions and their average confidence was 60%. The unpaid subjects answered correctly 43% of the questions and their average confidence was 53%. We first describe the results of the simple lotteries; the complementary (disjunctive) lotteries are discussed later.

The results are summarized by plotting the percentage of choices (C) that favor the judgment bet over the lottery as a function of judged probability (P). Before discussing the actual data, it is instructive to examine several constrasting predictions, implied by five alternative hypotheses, which are displayed in figure 1.

The upper panel of figure 1 displays the predictions of three hypotheses in which *C* is independent of *P*. According to expected utility theory, decision makers will be indifferent between betting on their judgment or betting on a chance lottery; hence *C* should equal 50% throughout. Ambiguity aversion implies that people will prefer to bet on a chance event whose probability is well defined rather than on their judged probability, which is inevitably vague; hence *C* should fall below 50% everywhere. The opposite hypothesis, called chance aversion, predicts that people will prefer to bet on their judgment rather than on a matched chance lottery; hence *C* should exceed 50% for all *P*. In contrast to the flat predictions displayed in the upper panel, the two hypothesis in the lower panel imply that *C* depends on *P*. The regression hypothesis states that the decision weights, which control choice, will be regressive relative to stated probabilities. Thus, *C* will be relatively high for small probabilities and relatively low for high probabilities. This prediction also follows from the theory proposed by Einhorn and Hogarth (1985), who put forth a particular process model based on mental simulation, adjustment, and anchoring. The predictions of this model, however, coincide with the regression hypothesis.

Finally, the competence hypothesis implies that people will tend to bet on their judgment when they feel knowledgeable and on the chance lottery when they feel ignorant.

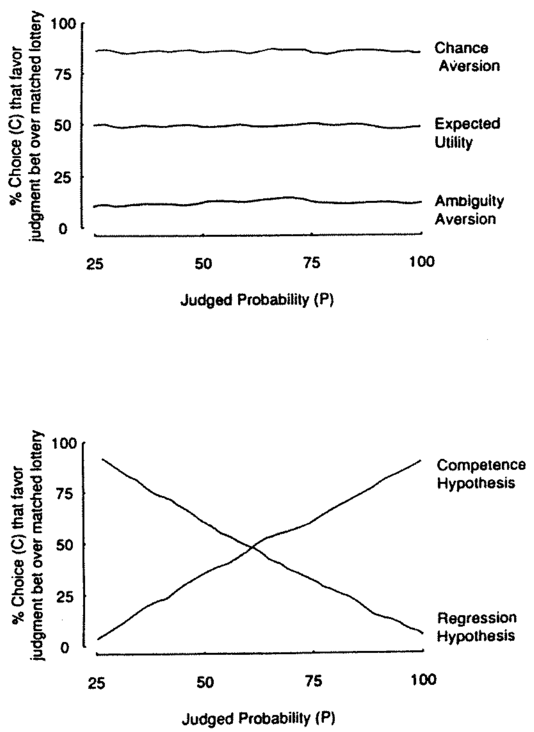

Figure 1. Five contrasting predictions of the results of an uncertainty preference experiment.

Because higher stated probability generally entails higher knowledge, C will be an increasing function of P except at 100% where the chance lottery amounts to a sure thing.

The results of the experiment are summarized in table 1 and figure 2. Table 1 presents, for three different ranges of P, the percentage of paid and nonpaid subjects who bet on their answers rather than on the matched lottery. Recall that each question had four possible answers, so the lowest confidence level is 25%. Figure 2 displays the overall percentage of choices C that favored the judgment bet over the lottery as a function 2 of judged probability P. The graph shows that subjects chose the lottery when P was low or moderate (below 65%) and that they chose to bet on their answers when P was high. The pattern of results was the same for the paid and for the nonpaid subjects, but the effect was slightly stronger for the latter group. These results confirm the prediction of the competence hypothesis and reject the four alternative accounts, notably the ambiguity aversion hypothesis implied by second-order probability models (e.g., Gärdenfors and Sahlin, 1982), and the regression hypothesis implied by the model of Einhorn and Hogarth (1985).

To obtain a statistical test of the competence hypothesis, we computed, separately for each subject, the binary correlation coefficient (ϕ) between choice (judgment bet vs. lottery) and judged probability (above median vs. below median). The median judgment was .65. Seventy-two percent of the subjects yielded positive coefficients, and the average ϕ was .30, ($t(54) = 4.3, p < .01$). To investigate the robustness of the observed pattern, we replicated the experiment with one major change. Instead of constructing chance lotteries whose probabilities matched the values stated by the subjects, we constructed lotteries in which the probability of winning was either 6% higher or 6% lower than the subjects' judged probability. For high-knowledge questions ($P \geq 75\%$), the majority of responses (70%) favored the judgment bet over the lottery even when the lottery offered a (6%) higher probability of winning. Similarly, for low-confidence questions ($P \geq 50\%$) the majority of responses (52%) favored the lottery over the judgment bet even when the former offered a lower (6%) probability of winning.

Figure 3 presents the calibration curve for the data of experiment 1. The figure shows that, on the whole, people are reasonably well calibrated for low probability but exhibit substantial overconfidence for high probability. The preference for the judgment bet over the lottery for high probability, therefore, cannot be justified on an actuarial basis.

The analysis of the complementary bets, where subjects were asked in effect to bet that their chosen answer was incorrect, revealed a very different pattern. Across subjects, the judgment bet was favored 40.5% of the time, indicating a statistically significant preference for the chance lottery ($t(54) = 3.8 p < .01$). Furthermore, we found no systematic

Table 1. Percentage of paid and nonpaid subjects who preferred the judgment bet over the lottery for low, medium, and high P (the number of observations are given in parentheses)

	$25 \leq P \leq 50$	$50 < P < 75$	$75 \leq P \leq 100$
Paid	29	42	55
	(278)	(174)	(168)
Nonpaid	22	43	69
	(394)	(188)	(140)

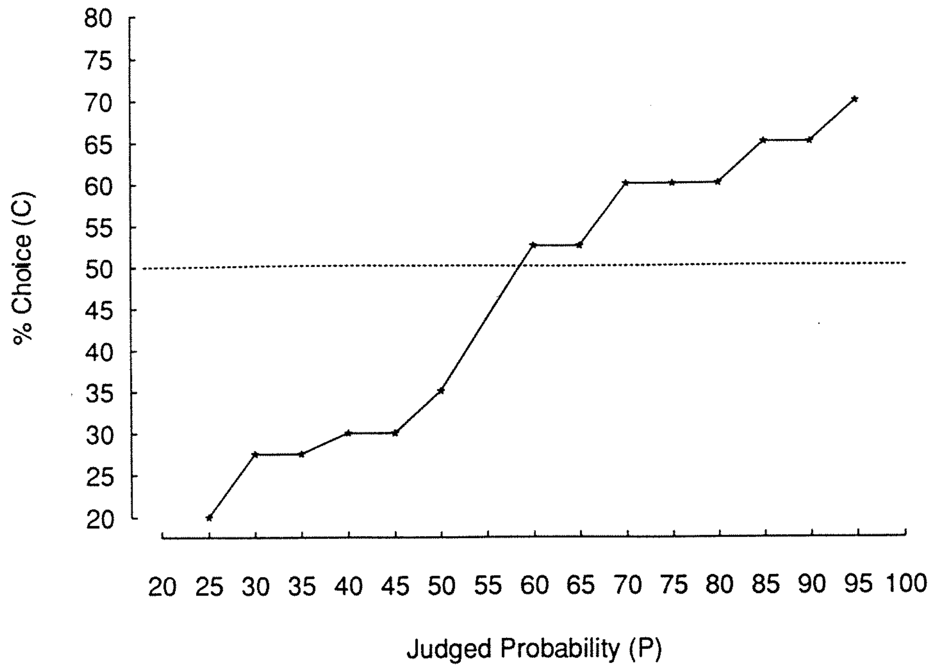

Figure 2. Percentage of choice (C) that favor a judgment bet over a matched lottery as a function of judged probability (P) in experiment 1.

relation between C and P, in marked contrast to the monotonic relation displayed in figure 2. In accord with our attributional account, this result suggests that people prefer to bet on their beliefs rather than against them. These data, however, may also be explained by the hypothesis that people prefer to bet on simple rather than on disjunctive hypotheses.

1.2. Experiment 2: Football and politics

Our next experiment differs from the previous one in three respects. First, it concerns the prediction of real-world future events rather than the assessment of general knowledge. Second, it deals with binary events so that the lowest level of confidence is .5 rather than .25 as in the previous experiment. Third, in addition to judgments of probability, subjects also rated their level of knowledge for each prediction.

A group of 20 students predicted the outcomes of 14 football games each week for five consecutive weeks. For each game, subjects selected the team that they thought would win the game and assessed the probability of their chosen team winning. The subjects also assessed, on a five-point scale, their knowledge about each game. Following the rating, subjects were asked whether they preferred to bet on the team they chose or on a matched chance lottery. The results summarized in figure 4 confirm the previous finding.

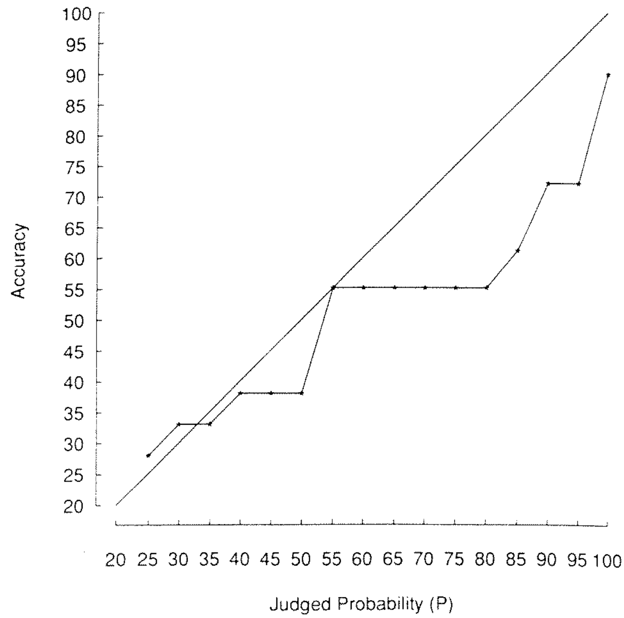

Figure 3. Calibration curve for experiment 1.

For both high and low knowledge (defined by a median split on the knowledge rating scale), C was an increasing function of P. Moreover, C was greater for high knowledge than for low knowledge at any $P > .5$. Only 5% of the subjects produced negative correlations between C and P, and the average ϕ coefficient was .33, ($t(77) = 8.7, p < .01$).

We next took the competence hypothesis to the floor of the Republican National Convention in New Orleans during August of 1988. The participants were volunteer workers at the convention. They were given a one-page questionnaire that contained instructions and an answer sheet. Thirteen states were selected to represent a cross section of different geographical areas as well to include the most important states in terms of electoral votes. The participants ($N = 100$) rates the probability of Bush carrying each of the 13 states in the November 1988 election on a scale from 0 (Bush is certain to lose) to 100 (Bush is certain to win). As in the football experiment, the participants rated their knowledge of each state on a five-point scale and indicated whether they would rather bet on their prediction or on a chance lottery. The results, summarized in figure 5, show that C increased with P for both levels of knowledge, and that C was greater for high knowledge than for low knowledge at all levels of P. When asked about their home state, 70% of the participants selected the judgment bet over the lottery. Only 5% of the subjects yielded negative correlations between C and P, and the average ϕ coefficient was .42, ($t(99) = 13.4, p < .01$).

The results displayed in figures 4 and 5 support the competence hypothesis in the prediction of real-world events: in both tasks C increases with P, as in experiment 1. In that study, however, probability and knowledge were perfectly correlated; hence the choice–judgment discrepancy could be attributed to a distortion of the probability scale

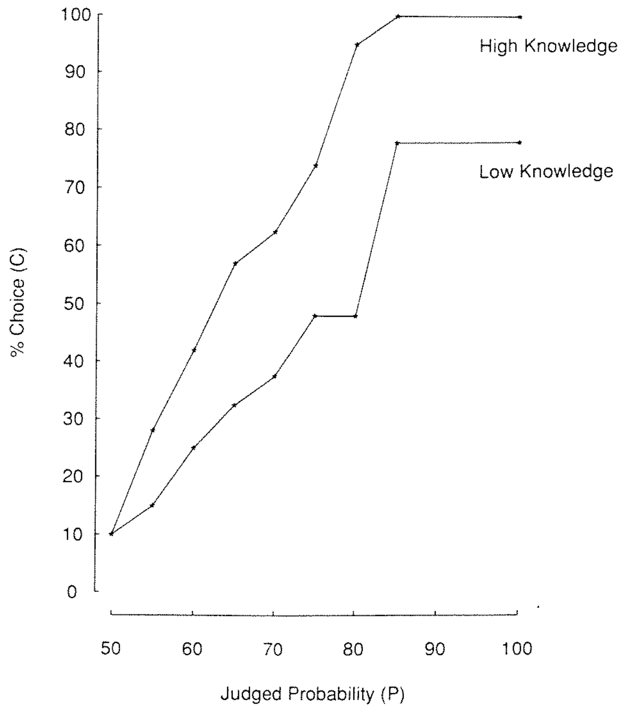

Figure 4. Percentage of choices (*C*) that favor a judgment bet over a matched lottery as a function of judged probability (*P*), for high- and low-knowledge items in the football prediction task (experiment 2).

in the judgement task. This explanation does not apply to the results of the present experiment, which exhibits an independent effect of rated knowledge. As seen in figures 4 and 5, the preference for the judgment bet over the chance lottery is greater for high-knowledge items than for low-knowledge items for all levels of judged probability. It is noteworthy that the strategy of betting on judgment was less successful than the strategy of betting on chance in both data sets. The former strategy yielded hit rates of 64% and 78% for football and election, respectively, whereas the latter strategy yielded hit rates of 73% and 80%. The observed tendency to select the judgment bet, therefore, does not yield better performance.

1.3. Experiment 3: Long shots

The preceding experiments show that people often prefer to bet on their judgment than on a matched chance event, even though the former is more ambiguouis than the latter. This effect, summarized in figures 2, 4 and 5, was observed at the high end of the probability scale. These data could perhaps be explained by the simple hypothesis that people prefer the judgment bet when the probability of winning exceeds .5 and the chance lottery when the probability of winning is below .5. To test this hypothesis, we

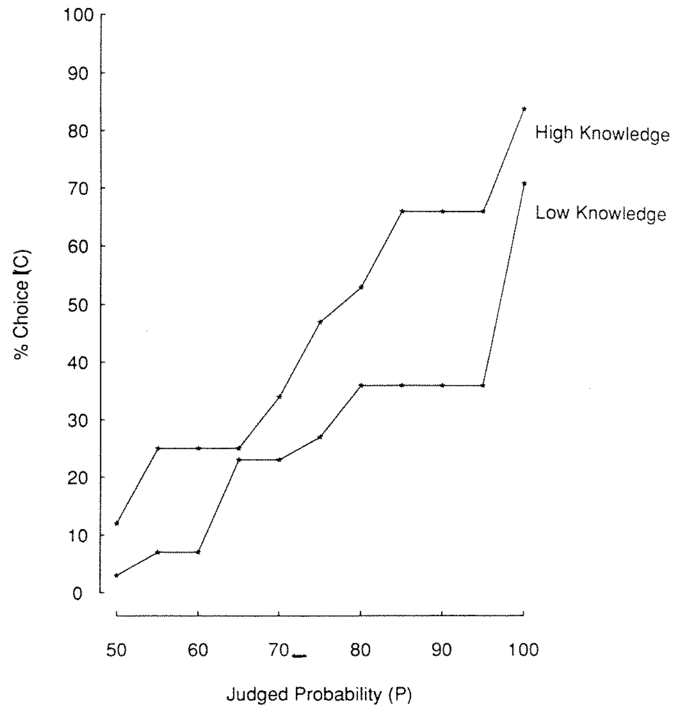

Figure 5. Percentage of choices (C) that favor a judgment bet over a matched lottery as a function of judged probability (P), for high- and low-knowledge items in experiment 2 (election data).

sought high-knowledge items in which the probability of winning is low, so the subject's best guess is unlikely to be true. In this case, the above hypothesis implies a preference for the chance lottery, whereas the competence hypothesis implies a preference for the judgment bet. These predictions are tested in the following experiment.

One hundred and eight students were presented with open-ended questions about 12 future events (e.g., what movie will win this year's Oscar for best picture? What football team will win the next Super Bowl? In what class next quarter will you have the highest grade?). They were asked to answer each question, to estimate the chances that their guess will turn out to be correct, and to indicate whether they have high or low knowledge of the relevant domain. The use of open-ended questions eliminates the lower bound of 50% imposed by the use of dichotomous predictions in the previous experiment. After the subjects completed these tasks, they were asked to consider, separately for each question, whether they would rather bet on their prediction or on a matched chance lottery.

On average, the subjects answered 10 out of 12 questions. Table 2 presents the percentage (C) of responses that favor the judgment bet over the chance lottery for high- and low-knowledge items, and for judged probabilities below or above .5. The number of responses in each cell is given in parentheses. The results show that, for high-knowledge

Table 2. Percentage of choices (C) that favor a judgment bet over a matched lottery for high- and low-rated knowledge and for judged probability below and above .5 (the number of responses are given in parentheses)

	Judged probability	
Rated knowledge	$P < .5$	$P \geq .5$
Low	36 (593)	58 (128)
High	61 (151)	69 (276)

items, the judgment bet was preferred over the chance lottery regardless of whether P was above or below one half ($p < .01$ in both cases), as implied by the competence hypothesis. Indeed, the discrepancy between the low- and high-knowledge conditions was greater for $P < .5$ than for $P \geq .5$. Evidently, people prefer to bet on their high-knowledge predictions even when the predictions are unlikely to be correct.

1.4. Experiment 4: Expert prediction

In the preceding experiments, we used the subjects' ratings of specific items to define high and know knowledge. In this experiment, we manipulate knowledge or competence by sorting subjects according to their expertise. To this end, we asked 110 students in an introductory psychology class to assess their knowledge of politics and of football on a nine-point scale. All subjects who rated their knowledge of the two areas on opposite sides of the midpoint were asked to take part in the experiment. Twenty-five subjects met this criterion, and all but two agreed to particpate. They included 12 political "experts" and 11 football "experts" defined by their strong area. To induce the subjects to give careful responses, we gave them detailed instructions including a discussion of calibration, and we employed the Brier scoring rule (see, e.g., Lichtenstein et al., 1982) designed to motivate subjects to give their best estimates. Subjects earned about \$10, on average.

The experiment consisted of two sessions. In the first session, each subject made predictions for a set of 40 future events (20 political events and 20 football games). All the events were resolved within five weeks of the date of the initial session. The political events concerned the winner of the various states in the 1988 presidential election. The 20 football games included 10 professional and 10 college games. For each contest (politics or football), subjects chose a winner by circling the name of one of the contestants, and then assessed the probability that their prediction would come true (on a scale from 50% to 100%).

Using the results of the first session, 20 triples of bets were constructed for each participant. Each triple included three matched bets with the same probability of winning generated by 1) a chance device, 2) the subject's prediction in his or her strong category, 3) the subject's prediction in his or her weak category. Obviously, some events appeared in more than one triple. In the second session, subjects ranked each of the 20 triples of bets. The chance bets were defined as in experiment 1 with reference to a box containing 100 numbered chips. Subjects were told that they would actually play their choices in each of the triples. To encourage careful ranking, subjects were told that they would play 80% of their first choices and 20% of their second choices.

The data are summarized in table 3 and in figure 6, which plots the attractiveness of the three types of bets (mean rank order) against judged probability. The results show a clear preference for betting on the strong category. Across all triples, the mean ranks were 1.64 for the strong category, 2.12 for the chance lottery, and 2.23 for the weak category. The difference among the ranks is highly significant ($p < .001$) by the Wilcoxen rank sum test. In accord with the competence hypothesis, people prefer to bet on their judgment in their area of competence, but prefer to bet on chance in an area in which they are not well informed. As expected, the lottery became more popular than the high-knowledge bet only at 100%. This pattern of result is inconsistent with an account based on ambiguity or second-order probabilities because both the high-knowledge and the low-knowledge bets are based on vague judgmental probabilities whereas the chance lotteries have clear probabilities. Ambiguity aversion could explain why low-knowledge bets are less attractive than either the high-knowledge bet or the chance bet, but it cannot explain the major finding of this experiment that the vague high-knowledge bets are preferred to the clear chance bets.

A noteworthy feature of figure 6, which distinguishes it from the previous graphs, is that preferences are essentially independent of P. Evidently, the competence effect is fully captured in this case by the contrast between the categories; hence the added knowledge implied by the judged probability has little or no effect on the choice among the bets.

Figure 7 presents the average calibration curves for experiment 4, separately for the high- and low-knowledge categories. These graphs show that judgments were generally overconfident: subjects' confidence exceeded their hit rate. Furthermore, the overconfidence was more pronounced in the high-knowledge category than in the low-knowledge category. As a consequence, the ordering of bets did not mirror judgmental accuracy. Summing across all triples, betting on the chance lottery would win 69% of the time, betting on the novice category would win 64% of the time, and betting on the expert category would win only 60% of the time. By betting on the expert category therefore the subjects are losing, in effect, 15% of their expected earnings.

The preference for knowledge over chance is observed not only for judgments of probability for categorical events (win, loss), but also for probability distributions over numerical variables. Subjects ($N = 93$) were given an opportunity to set 80% confidence intervals for a variety of quantities (e.g., average SAT score for entering freshmen at

Table 3. Ranking data for expert study

	Rank			
Type of bet	1st	2nd	3rd	Mean rank
High-knowledge	192	85	68	1.64
Chance	74	155	116	2.12
Low-knowledge	79	105	161	2.23

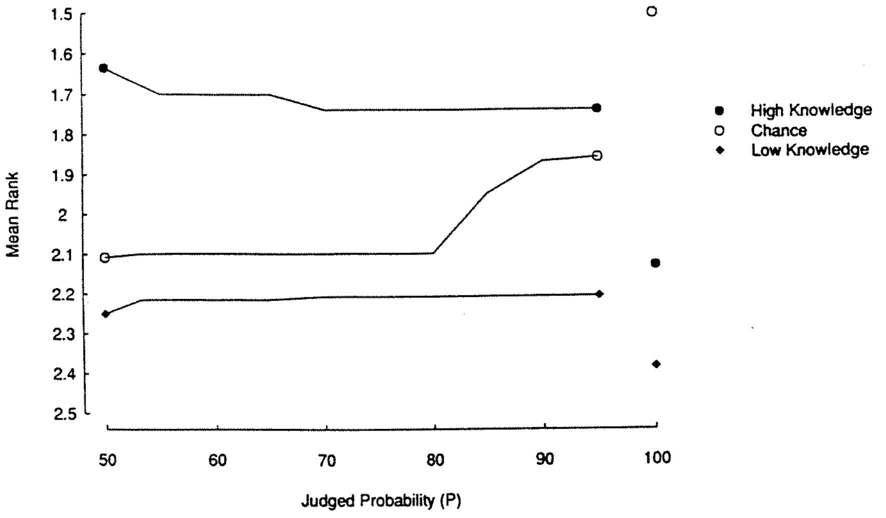

Figure 6. Ranking data for high-knowledge, low-knowledge, and chance bets as a function of *P* in experiment 4.

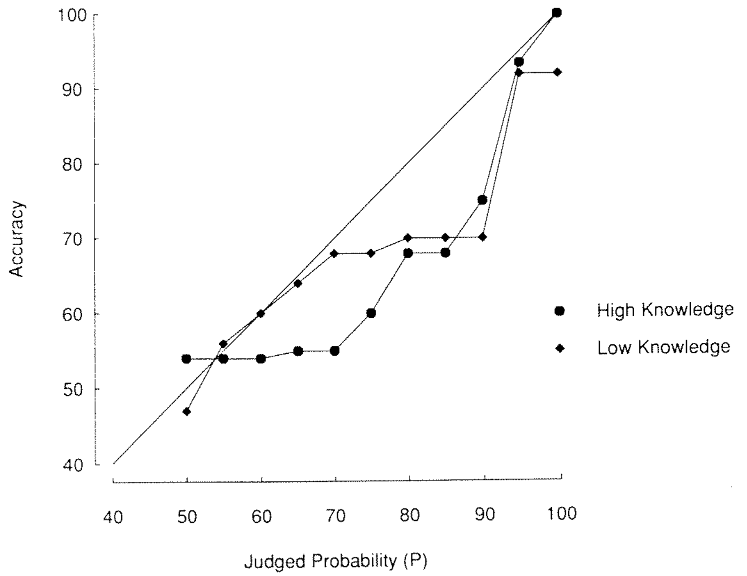

Figure 7. Calibration curves for high- and low-knowledge categories in experiment 4.

Stanford; driving distance from San Francisco to Los Angeles). After setting confidence intervals, subjects were given the opportunity to choose between 1) betting that their confidence interval contained the true value, or 2) an 80% lottery. Subjects preferred betting on the confidence interval in the majority of cases, although this strategy paid off only 69% of the time because the confidence intervals they set were generally too narrow. Again, subjects paid a premium of nearly 15% to bet on their judgment.

1.5. *Experiment 5: Complementary bets*

The preceding experiments rely on judgments of probabillity to match the chance lottery and the judgment bet. To control for possible biases in the judgment process, our last test of the competence hypothesis is based on a pricing task that does not involve probability judgment. This experiment also provides an estimate of the premium that subjects are paying in order to bet on high-knowledge items.

Sixty-eight students were instructed to state their cash equivalent (reservation price) for each of 12 bets. They were told that one pair of bets would be chosen and a few students, selected at random, would play the bet for which they stated the higher cash equivalent. (For a discussion of this payoff scheme, see Tversky, Slovic and Kahneman, 1990.) All bets in this experiment offered a prize of \$15 if a given proposition were true, and nothing otherwise. Complementary propositions were presented to different subjects. For example, half the subjects were asked to price a bet that paid \$15 if the air distance between New York and San Francisco is more 2500 miles, and nothing otherwise. The half of the subjects were asked to price the complementary bet that paid \$15 if the air distance between New York and San Francisco is less than 2500 miles, and nothing otherwise.

To investigate uncertainty preferences, we paired high-knowledge and low-knowledge propositions. For example, we assumed that the subjects know more about the air distance between New York and San Francisco than about the air distance between Beijing and Bangkok. We also assumed that our respondents (Stanford students) know more about the percentage of undergraduate students who receive on-campus housing at Stanford than at the University of Nevada, Las Vegas. As before, we refer to these propositions as high-knowledge and low-knowledge items, respectively. Note that the selection of the stated value of the uncertain quantity (e.g., air distance, percentage of students) controls a subject's confidence in the validity of the proposition in question, independent of his or her general knowledge about the subject matter. Twelve pairs of complementary propositions were constructed, and each subject evaluated one of the four bets defined by each pair. In the air-distance problem, for example, the four propositions were d(SF,NY) > 2500, d(SF,NY) < 2500, d(Be,Ba) > 3000, d(Be,Ba) < 3000, where d(SF,NY) and d(Be,Ba) denote, respectively, the distances between San Francisco and New York and between Beijing and Bangkok.

Note that according to expected value, the average cash equivalent for each pair of complementary bets should be \$7.50. Summing across all 12 pairs of complementary bets, subjects paid on average \$7.12 for the high-knowledge bets and only \$5.96 for the low-knowledge bets ($p < .01$). Thus, people were paying, in effect, a competence premium of nearly 20% in order to bet on the more familiar propositions. Furthermore, the average price for the (complementary) high-knowledge bets was greater than that for the low-knowledge bets in 11 out of 12 problems. For comparison, the average cash equivalent for a coin toss to win \$15 was \$7. In accord with our previous findings, the chance lottery is valued above the low-knowledge bets but not above the high-knowledge bets.

We next test the competence hypothesis against expected utility theory. Let H and $\overline{H}$ denote two complimentary high-knowledge propositions, and let L and $\overline{L}$ denote the corresponding low-knowledge propositions. Suppose a decision maker prefers betting on

H over L and on $\overline{H}$ over $\overline{L}$. This pattern is inconsistent with expected utility theory because it implies that $P(H) > P(L)$ and $P(\overline{H}) > P(\overline{L})$, contrary to the additivity assumption $P(H) + P(\overline{H}) = P(L) + P(\overline{L}) = 1$. If, on the other hand, high-knowledge bets are preferred to low-knowledge bets, such a pattern is likely to arise. Because the four propositions $(H, \overline{H}, L, \overline{L})$ were evaluated by four different groups of subjects, we employ a between-subject test of additivity. Let $M(H_i)$ be the median price for the high-knowledge proposition H_i, etc. The responses in problem i violate additivity, in the direction implied by the competence hypothesis, whenever $M(H_i) > M(L_i)$ and $M(\overline{H}_i) \geq M(\overline{L}_i)$.

Five of the 12 pairs of problems exhibited this pattern indicating a preference for the high-knowledge bets, and none of the pairs exhibited the opposite pattern. For example, the median price for betting on the proposition "more than 85% of undergraduates at Stanford receive on-campus housing" was \$7.50, and the median cash equivalent for betting on the complementary proposition was \$10. In contrast, the median cash equivalent for betting on the proposition "more than 70% of undergraduates at UNLV receive on-campus housing" was \$3, and the median value for the complementary bet was \$7. The majority of respondents, therefore, were willing to pay more to bet on either side of a high-knowledge item than on either side of a low-knowledge item.

The preceding analysis, based on medians, can be extended as follows. For each pair of propositions (H_i, L_i), we computed the proportion of comparisons in which the cash equivalent of H_i exceeded the cash equivalent of L_i, denoted $P(H_i > L_i)$. We also computed $P(\overline{H}_i > \overline{L}_i)$ for the complementary propositions. All ties were excluded. Under expected utility theory,

$$P(H_i > L_i) + P(\overline{H}_i > \overline{L}_i) = P(L_i > H_i) + P(\overline{L}_i > \overline{H}_i) = 1,$$

because the additivity of probability implies that for every comparison that favors H_i over L_i, there should be another comparison that favors $\overline{L}_i$ over $\overline{H}_i$. On the other hand, if people prefer the high-knowledge bets, as implied by the competence hypothesis, we expect

$$P(H_i > L_i) + P(\overline{H}_i > \overline{L}_i) > P(L_i > H_i) + P(\overline{L}_i > \overline{H}_i).$$

Among the 12 pairs of complementary propositions, the above inequality was satisfied in 10 cases, the opposite inequality was satisfied in one case, and equality was observed in one case, indicating a significant violation of additivity in the direction implied by the competence hypothesis ($p < .01$ by sign test). These findings confirm the competence hypothesis in a test that does not rely on judgments of probability or on a comparison of a judgment bet to a matched lottery. Hence, the present results cannot be attributed to a bias in the judgment process or in the matching of high- and low-knowledge items.

2. Discussion

The experiments reported in this article establish a consistent and pervasive discrepancy between judgments of probability and choice between bets. Experiment 1 demonstrates that the preference for the knowledge bet over the chance lottery increases with judged confidence. Experiments 2 and 3 replicate this finding for future real-world events, and demonstrate a knowledge effect independent of judged probability. In experiment 4, we sort subjects into their strong and weak areas and show that people like betting on their strong category and dislike betting on their weak category; the chance bet is intermediate between the two. This pattern cannot be explained by ambiguity or by second-order probability because chance is unambiguous, whereas judgmental probability is vague. Finally, experiment 5 confirms the prediction of the competence hypothesis in a pricing task that does not rely on probability matchings, and shows that people are paying a premium of nearly 20% for betting on high-knowledge items.

These observations are consistent with our attributional account, which holds that knowledge induces an asymmetry in the internal balance of credit and blame. Competence, we suggest, allows people to claim credit when they are right, and its absence exposes people to blame when they are wrong. As a consequence, people prefer the high-knowledge bet over the matched lottery, and they prefer the matched lottery over the low-knowledge bet. This account explains other instances of uncertainty preferences reported in the literature, notably the preference for clear over vague probabilities in a chance setup (Ellsberg, 1961), the preference to bet on the future over the past (Rothbart and Snyder, 1970, Brun and Teigen, 1989), the preference for skill over chance (Cohen and Hansel, 1959; Howell, 1971), and the enhancement of ambiguity aversion in the presence of knowledgeable others (Curley, Yates and Abrams, 1986). The robust finding that, in their area of competence, people prefer to bet on their (vague) beliefs over a matched chance event indicates that the impact of knowledge or competence outweighs the effect of vagueness.

In experiments 1-4 we used probability judgments to establish belief and choice data to establish preference. Furthermore, we have interpreted the choice-judgment discrepancy as a preference effect. In contrast, it could be argued that the choice-judgment discrepancy is attributable to a judgmental bias, namely underestimation of the probabilities of high-knowledge items and an overestimation of the probabilities of low-knowledge items. This interpretation, however, is not supported by the available evidence. First, it implies less overconfidence for high-knowledge than for low-knowledge items contrary to fact (see figure 7). Second, judgments of probability cannot be dismissed as inconsequential because in the presence of a scoring rule, such as the one used in experiment 4, these judgments represent another form of betting. Finally, a judgmental bias cannot explain the results of experiment 5, which demonstrates preferences for betting on high-knowledge items in a pricing task that does not involve probability judgment.

The distinction between preference and belief lies at the heart of Bayesian decision theory. The standard interpretation of this theory assumes that 1) the expressed beliefs (i.e., probabililty judgments) of an individual are consistent with an additive probability measure, 2) the preferences of an individual are consistent with the expectation principle, and hence give rise to a (subjective) probability measure derived from choice, and 3) the two measures of subjective probability—obtained from judgment and from choice—are consistent. Note that points 1 and 2 are logically independent. Allais' counterexample, for instance, violates 2 but not 1. Indeed, many authors have introduced nonadditive decision weights, derived from preferences, to accommodate the observed violations of the expectation principle (see, e.g., Kahneman and Tversky, 1979). These decision weights, however, need not reflect the decision maker's beliefs. A person may believe that the probability of drawing the ace of spades from a well-shuffled deck is 1/52, yet in betting on this event he or she may give it a higher weight. Similarly, Ellsberg's example does not prove that people regard the clear event as more probable than the corresponding vague event; it only shows that people prefer to bet on the clear event. Unfortunately, the term *subjective probability* has been used in the literature to describe decision weights derived from preference as well as direct expressions of belief. Under the standard interpretation of the Bayesian theory, the two concepts coincide. As we go beyond this theory, however, it is essential to distinguish between the two.

2.1. Manipulations of ambiguity

The distinction between belief and preference is particularly important for the interpretation of ambiguity effects. Several authors have concluded that, when the probability of winning is small or when the probability of losing is high, people prefer ambiguity to clarity (Curley and Yates, 1989; Einhorn and Hogarth, 1985; Hogarth and Kunreuther, 1989). However, this interpretation can be challenged because, as will be shown below, the data may reflect differences in belief rather than uncertainty preferences. In this section, we investigate the experimental procedures used to manipulate ambiguity and argue that they tend to confound ambiguity with perceived probability.

Perhaps the simplest procedure for manipulating ambiguity is to vary the decision maker's confidence in a given probability estimate. Hogarth and his collaborators have used two versions of this procedure. Einhorn and Hogarth (1985) presented the subject with a probability estimate, based on the "judgement of independent observers," and varied the degree of confidence attached to that estimate. Hogarth and Kunreuther (1989) "endowed" the subject with his or her "best estimate of the probability" of a given event, and manipulated ambiguity by varying the degree of confidence associated with this estimate. If we wish to interpret people's willingness to bet on these sorts of events as ambiguity seeking or ambiguity aversion, however, we must first verify that the manipulation of ambiguity did not affect the perceived probability of the events.

To investigate this question, we first replicated the manipulation of ambiguity used by Hogarth and Kunreuther (1989). One group of subjects ($N = 62$), called the high confidence group, received the following information:

> Imagine that you head a department in a large insurance company. The owner of a small business comes to you seeking insurance against a $100,000 loss which could result from claims concerning a defective product. You have considered the manufacturing process, the reliabilities of the machines used, and evidence contained in the business records. After considering the evidence available to you, your best estimate of the probability of a defective product is .01. Given the circumstances, you feel confident about the precision of this estimate. Naturally you will update your estimate as you think more about the situation or receive additional information.

A second group of subjects ($N = 64$), called the low-confidence group, received the same information, except that the phrase "you feel confident about the precision of this estimate" was replaced by "you experience considerable uncertainty about the precision of this estimate." All subjects were then asked:

> Do you expect that the new estimate will be (Check one):
> Above .01 ____________
> Below .01 ____________
> Exactly .01 ____________

The two groups were also asked to evaluate a second case in which the stated probability of a loss was .90. If the stated value (.01 or .90) is interpreted as the mean of the respective second-order probability distribution, then a subject's expectation for the updated estimate should coincide with the current "best estimate." Furthermore, if the manipulation of confidence affects ambiguity but not perceived probability, there should be no difference between the responses of the high-confidence and the low-confidence groups. The data presented in table 4, under the heading *Your probability,* clearly violate these assumptions. The distributions of responses in the low-confidence condition are considerably more skewed than the distributions in the high-confidence condition. Furthermore, the skewness is positive for .01 and negative for .90. Telling subjects that they "experience considerable uncertainty" about their best estimate produces a regressive shift: the expected probability of loss is above .01 in the first problem and below .90 in the second. The interaction between confidence (high–low) and direction (above–below) is statistically significant ($p < .01$).

We also replicated the procedure employed by Einhorn and Hogarth (1985) in which subjects were told that "independent observers have stated that the probability of a defective product is .01." Subjects ($N = 52$) in the high-confidence group were told that "you could feel confident about the estimate," whereas subjects ($N = 52$) in the low-confidence group were told that "you could experience considerable uncertainty about the estimate." Both groups were then asked whether their best guess of the probability of experiencing a loss is above .01, below .01, or exactly .01. The two groups also evaluated a

Table 4. Subjective assessments of stated probabilities of .01 and .90 under high-confidence and low-confidence instructions (the entries are the percentage of subjects who chose each of the three responses)

		Your probability		Others' estimate	
Stated value	Response	High confidence	Low confidence	High confidence	Low confidence
	Above .01	45	75	46	80
.01	Exactly .01	34	11	15	6
	Below .01	21	14	39	14
	Above .90	29	28	42	26
.90	Exactly .90	42	14	23	12
	Below .90	29	58	35	62

second case in which the probability of loss was .90. The results presented in table 4, under the heading *Others' estimate,* reveal the pattern observed above. In the high-confidence condition, the distributions of responses are fairly symmetric, but in the low-confidence condition the distributions exhibit positive skewness at .01 and negative skewness at .90. Again, the interaction between confidence (high–low) and direction (above–below) is the statistically significant ($p < .01$).

These results indicate that the manipulations of confidence influenced not only the ambiguity of the event in question but also its perceived probability: they increased the perceived probability of the highly unlikely event and decreased the perceived probability of the likely event. A regressive shift of this type is not at all unreasonable and can even be rationalized by a suitable prior distribution. As a consequence of the shift in probability, the bet on the vaguer estimate should be more attractive when the probability of loss is high (.90) and less attractive when the probability of loss is low (.01). This is exactly the pattern of preferences observed by Einhorn and Hogarth (1985) and by Hogarth and Kunreuther (1989), but it does not entail either ambiguity seeking or ambiguity aversion because the events differ in perceived probability, not only in ambiguity.

The results of table 4 and the findings of Hogarth and his collaborators can be explained by the hypothesis that subjects interpret the stated probability value as the median (or the mode) of a second-order probability distribution. If the second-order distributions associated with extreme probabilities are skewed towards .5, the mean is less extreme than the median, and the difference between them is greater when ambiguity is high than when it is low. Consequently, the mean of the second-order probability distribution, which controls choice in the Bayesian model, will be more regressive (i.e., closer to .5) under low confidence than under high confidence.

The potential confounding of ambiguity and degree of belief arises even when ambiguity is manipulated by information regarding a chance process. Unlike Ellsberg's comparison of the 50/50 box with the unknown box, where symmetry precludes a bias in one direction or another, similar manipulations of ambiguity in asymmetric problems could produce a regressive shift, as demonstrated in an unpublished study by Parayre and Kahneman.[3]

These investigators compared a clear event, defined by the proportion of red balls in a box, with a vague event defined by the range of balls of the designated color. For a vague event [.8, 1], subjects were informed that the proportion of red balls could be anywhere between .8 and 1, compared with .9 for the clear event. Table 5 presents both choice and judgment data for three probability levels: low, medium, and high. In accord with previous work, the choice data show that subjects preferred to bet on the vague event when the probability of winning was low and when the probability of losing was high, and they preferred to bet on the clear event in all other cases. The novel feature of the Parayre and Kahneman experiment is the use of a perceptual rating scale based on a judgment of length, which provides a nonnumerical assessment of probability. Using this scale, the investigators showed that the judged probabilities were regressive. That is, the vague low-probability event [0,.10] was judged as more probable than the clear event, .05, and the vague high-probability event [.8,1] was judged as less probable than the clear event, .9. For the medium probability, there was no significant difference in judgment between the vague event [0,1] and the clear event, .5. These results, like the data of table 4,

Table 5. (Based on Parayre and Kahneman). Percentage of subjects who favored the clear event and the vague event in judgment and in choice.

	Probability (win/lose)	Judgment N = 72	Choice: Win \$100 N = 58	Choice: Lose \$100 N = 58
Low	.05	28	12	66
	[0,.10]	47	74	22
Medium	.5	38	60	60
	[0,1]	22	26	21
High	.9	50	50	22
	[.8,1]	21	34	47

Note: The sum of the two values in each condition is less than 100%; the remaining responses expressed equivalence. In the choice task, the low probabilities were .075 and [0,.15]. N denotes sample size.

demonstrate that the preference for betting on the ambiguous event (observed at the low end for positive bets and at the high end for negative bets) could reflect a regressive shift in the perception of probability rather than a preference for ambiguity.

2.2. *Concluding remarks*

The findings regarding the effect of competence and the relation between preferences and beliefs challenge the standard interpretation of choice models that assumes independence of preference and belief. The results are also at variance with post-Bayesian models that invoke second-order beliefs to explain the effects of ambiguity or partial knowledge. Moreover, our results call into question the basic idea of defining beliefs in terms of preferences. If willingness to bet on an uncertain event depends on more than the perceived likelihood of that event and the confidence in that estimate, it is exceedingly difficult—if not impossible—to derive underlying beliefs from preferences between bets.

Besides challenging existing models, the competence hypothesis might help explain some puzzling aspects of decisions under uncertainty. It could shed light on the observation that many decision makers do not regard a calculated risk in their area of competence as a gamble (see, e.g., March and Shapira, 1987). It might also help explain why investors are sometimes willing to forego the advantage of diversification and concentrate on a small number of companies (Blume, Crockett, and Friend, 1974) with which they are presumably familiar. The implications of the competence hypothesis to decision making at large are left to be explored.

3. Notes

1. We use the term *competence* in a broad sense that includes skill, as well as knowledge or understanding.

2. In this and all subsequent figures, we plot the isotone regression of C on P—that is, the best-fitting monotone function in the least squares sense (see Barlow, Bartholomew, Bremner and Brunk, 1972).

3. We are grateful to Parayre and Kahneman for providing us with these data.

References

Anscombe, F.J., and R.J. Aumann, (1963). "A Definition of Subjective Probability," *Annals of Mathematical Statistics,* 34, 199-205.

Barlow, Richard E., et al. (1972). *Statistical Inference Under Order Restrictions: The Theory and Application of Isotonic Regression.* New York: J. Wiley.

Blume, Marshall E., Jean Crockett, and Irwin Friend. (1974). "Stock Ownership in the United States: Characteristics and Trends," *Survey of Current Business.*

Brun, Wibecke, and Karl Teigen. (1990). "Prediction and Postdiction Preferences in Guessing," *Journal of Behavioral Decision Making* 3, 17-28.

Budescu, David, Shalva Weinberg, and Thomas Wallsten. (1988). "Decisions Based on Numerically and Verbally Expressed Uncertainties," *Journal of Experimental Psychology: Human Perception and Performance* 14 (2), 281-294.

Cohen, John, and Mark Hansel. (1959). "Preferences for Different Combinations of Chance and Skill in Gambling," *Nature* 183, 841-843.

Curley, Shawn, and J. Frank Yates. (1989). "An Empirical Evaluation of Descriptive Models of Ambiguity Reactions in Choice Situations," *Journal of Mathematical Psychology* 33, 397-427.

Curley, Shawn, J. Frank Yates, and Richard Abrams. (1986). "Psychological Sources of Ambiguity Avoidance," *Organizational Behavior and Human Decision Processes* 38, 230-256.

Einhorn, Hillel, and Robin Hogarth. (1985). "Ambiguity and Uncertainty in Probabilistic Inference," *Psychological Review* 93, 433-461.

Ellsberg, Daniel. (1961). "Risk, Ambiguity, and the Savage Axioms," *Quarterly Journal of Economics* 75, 643-669.

Ellsberg, Daniel. (1963). "Risk, Ambiguity, and the Savage Axioms: Reply," *Quarterly Journal of Economics* 77, 336, 342

Fellner, William. (1961). "Distortion of Subjective Probabilities as a Reaction to Uncertainty," *Quarterly Journal of Economics* 75, 670-689.

Fishburn, Peter. (1988). *Nonlinear Preference and Utility Theory.* Baltimore, MD: The Johns Hopkins University Press.

Frisch, Deborah, and Jonathan Baron. (1988). "Ambiguity and Rationality," *Journal of Behavioral Decision Making* 1, 149-157.

Gärdenfors, Peter, and Nils-Eric Sahlin. (1982). "Unreliable Probabilities, Risk Taking, and Decision Making," *Synthese* 53 (3), 361-386.

Hogarth, Robin, and Howard Kunreuther. (1988). *Pricing Insurance and Warranties: Ambiguity and Correlated Risks.* Unpublished manuscript, University of Chicago and University of Pennsylvania.

Hogarth, Robin, and Howard Kunreuther. (1989). "Risk, Ambiguity, and Insurance," *Journal of Risk and Uncertainty* 2, 5-35.

Howell, William. (1971). "Uncertainty from Internal and External Sources: A Clear Case of Overconfidence," *Journal of Experimental Psychology* 89 (2), 240-243.

Kahneman, Daniel, and Amos Tversky. (1979). "Prospect Theory: An Analysis of Decision under Risk," *Econometrica* 47, 263–291.

Lichtenstein, Sarah, Baruch Fischhoff, and Lawrence Phillips. (1982). "Calibration of Probabilities: The State of the Art to 1980." In D. Kahneman, P. Slovic and A. Tversky (eds.), *Judgment under Uncertainty: Heuristics and Biases.* New York: Cambridge University Press.

March, James, and Zur Shapira. (1987). "Managerial Perspectives on Risk and Risk Taking," *Management Science* 33 (11), 1404–1418.

Raiffa, Howard. (1961). "Risk, Ambiguity, and the Savage Axioms: Comment," *Quarterly Journal of Economics* 75, 690–694.

Ramsey, Frank. (1931). "Truth and Probability." In F.P. Ramsey, *The Foundations of Mathematics and Other Logical Essays.* New York: Harcourt, Brace and Co.

Rothbart, Myron, and Mark Snyder. (1970). "Confidence in the Prediction and Postdiction of an Uncertain Outcome," *Canadian Journal of Behavioral Science* 2 (1), 38–43.

Savage, Leonard. (1954). *The Foundations of Statistics.* New York: Wiley.

Schmeidler, David. (1989). "Subjective Probability and Expected Utility Without Additivity," *Econometrica* 57 (3), 571–587.

Skyrm, Brian. (1980). "Higher Order Degrees of Belief." In D.H. Mellor (ed.), *Prospects for Pragmatism. Essays in Memory of F.P. Ramsey.* Cambridge: Cambridge University Press, pp. 109–137.

Tversky, Amos, Paul Slovic, and Daniel Kahnemann. (1990). "The Causes of Preference Reversal," *American Economic Review* 80, 204–217.