



Receiving other people's advice: Influence and benefit[☆]

Ilan Yaniv

Department of Psychology, Hebrew University of Jerusalem, Mt. Scopus, Jerusalem 91905, Israel

Abstract

Seeking advice is a basic practice in making real life decisions. Until recently, however, little attention has been given to it in either empirical studies or theories of decision making. The studies reported here investigate the influence of advice on judgment and the consequences of advice use for judgment accuracy. Respondents were asked to provide final judgments on the basis of their initial opinions and advice presented to them. The respondents' weighting policies were inferred. Analysis of these policies show that (a) the respondents tended to place a higher weight on their own opinion than on the advisor's opinion (the self/other effect); (b) more knowledgeable individuals discounted the advice more; (c) the weight of advice decreased as its distance from the initial opinion increased; and (d) the use of advice improved accuracy significantly, though not optimally. A theoretical framework is introduced which draws in part on insights from the study of attitude change to explain the influence of advice. Finally the usefulness of advice for improving judgment accuracy is considered.

© 2003 Elsevier Inc. All rights reserved.

We are usually convinced more easily by reasons we have found ourselves than by those which have occurred to others. – Blaise Pascal

The use of advice is a fundamental practice in making real-life decisions, whether as basic as finding directions in an unfamiliar environment or as complex as those involving legal or medical issues. However, until recently the use of advice has been given little consideration in either empirical studies or theories of decision making (Harvey & Fischer, 1997; Jonas & Frey, 2003; Jungermann, 1997; Sniezek & Buckley, 1995; Yaniv & Kleinberger, 2000). Advice seeking is important because real decision problems generally do not come as completely packaged, self-contained "textbook problems." Hence people engage in interactive social and cognitive processes of giving and taking advice to enhance their representation of a decision problem (Yates, Price, Lee, & Ramirez, 1996; Zarnoth & Sniezek, 1997). In particular, they solicit opinions from worthy advisors, assess their merit, and then combine them. An advisor might

fill in missing information, help assess the values of alternative options, or serve as a "sounding board." In sum, it appears that the use of advice plays a far greater role in the *practice* of real life decision making than it has had in decision research.

A major motivation for seeking advice is the need to improve judgment accuracy and the expectation that advice will help. An abundance of studies have shown that combining multiple sources of information improves estimation in the long run, in a variety of domains ranging from perceptual judgment to business forecasting (e.g., Armstrong, 2001; Sorkin, Hayes, & West, 2001; Yaniv, 1997). Aside from accuracy, there are also social reasons for seeking advice, which we consider only briefly here. Accountants performing complex audit tasks tend to solicit advice for self-presentational reasons and to increase the justification for their decisions (Kennedy, Kleinmuntz, & Peecher, 1997). Indeed, seeking advice implies sharing with others the responsibility for the outcome of a decision (Harvey & Fischer, 1997). One might argue, however, that even self-presentational reasons for seeking advice are rooted in the belief on the part of the individual or the organization that consulting someone else's opinion could improve one's final decision.

Whereas advising per se has received little attention in the study of decision making, several important lines of

[☆] This research was supported by Grant No. 822/00 from the Israel Science Foundation. The author is a member of the Department of Psychology and of the Center for the Study of Rationality, Hebrew University of Jerusalem.

E-mail address: ilan.yaniv@huji.ac.il.

research form the basis for the present investigation. These include theories in the following domains: (a) processes of attitude change, belief revision and perseverance (Zimbardo & Leippe, 1991), (b) the literature on combining expert opinions and linear models of judgment (Armstrong, 2001; Blattberg & Hoch, 1990), (c) models of information integration (Anderson, 1968), and (d) interactive group judgment (Davis et al., 1997; Sniezek & Henry, 1989). Research in these areas highlights the processes by which information is combined and opinions are revised.

The focus of the present research is on two aspects of advice seeking—how the advice is used and whether there is a resulting gain in accuracy. In these studies we consider perhaps the simplest form of advice use, namely getting a piece of information (numerical estimate) from an outside party and using it to update one's own view. As simple as it is, numerical advice has an important function in individual as well as organizational decisions. Physicians, weather forecasters, genetic consultants and lawyers, just to name a few, are all in the business of communicating their forecasts and uncertain estimates to others facing decisions. In a different vein, the use of numerical estimates has certain methodological advantages, primarily the ability to measure straightforwardly respondents' *weighting policies* and *accuracy gains*.

Policies for using advice

A basic dilemma in using advice involves the amount of weight to place on others' opinions. Receiving advice often exposes decision makers to a potential conflict between their initial opinions and the advice. Consider a manager who believes that a certain new product is likely to gain success and is thus worthy of further development. The manager then receives a lukewarm expert opinion of her idea. How might she revise her opinion? The key question in many practical situations is to decide just how much weight ought to be assigned to a particular piece of advice. In particular, a decision maker's weighting policy might entail completely ignoring the other opinion, some adjustment of one's own opinion towards the other, or complete adoption of the other opinion.

The studies presented here investigate how people weight others' opinions and how this weighting policy changes as a function of knowledge and of the distance of the advice from the decision maker's own opinion. Finally, the consequences of such policies for judgment accuracy are considered.

In order to develop hypotheses about the policies that decision makers use for integrating advice, I made use of an analogy between advice *use* and attitude *change*. The process of weighting advice in judgment may resemble

the processes underlying opinion change as a function of communication. To be sure, research in these two areas arises from different conceptual perspectives. Studies of judgment typically ask how good a person's judgment is in terms of its accuracy or coherence. Studies of attitudes are typically focused on the valence (e.g., positive vs negative) and strength of the person's attitude, with the goal of understanding what affects them (Ajzen, 2001). Moreover, in attitude change the main perspective is that of the communicator, who seeks to influence or persuade target recipients (Zimbardo & Leippe, 1991). In advice seeking, the recipient often initiates the process in attempt to improve the quality of her judgment. The goal of influence promotion is manipulative—that is, bringing about change in some preferred direction—whereas a major goal in seeking advice is improving decision quality.

Despite these differences, it is not inconceivable that advice use and attitude change share certain commonalities. In both cases one's initial opinion is integrated with that of someone else, be it a communicator's influential message or an advisor's opinion. I pursue the merits (as well as the limits) of this analogy in subsequent sections and the final discussion. Drawing on this analogy, I outline two hypotheses which involve the mechanisms that underlie discounting and the effect of distance. Both reflect the manner in which judges resolve the conflict between their initial opinions and the advice. I also consider the consequences of advice use for accuracy.

The self/other effect: Discounting the weight of advice

Previous work on the use of advice in decision making suggests a self/other effect whereby individuals tend to discount advice and favor their own opinion. In a judgmental estimation task (Yaniv & Kleinberger, 2000) respondents formed a final opinion on the basis of their initial opinion and a piece of advice. Rather than using equal weighting, respondents tended to place a higher weight on their own opinion than on the advisor's opinion. Even though the decision makers were sensitive to the quality of the advice (good vs poor), they tended to discount both good and poor advice. In a cue-learning study by Harvey and Fischer (1997), respondents shifted their estimates about 20–30% towards the advisor's estimates. Lim and O'Connor (1995) found that, in combining their prior personal forecasts and advisory (statistical) forecasts, judges weighted their own forecasts more heavily than the statistical forecasts.

I suggest that these discounting phenomena result from the nature of the support the judge can recruit for her own opinion versus the advice. In particular, the self/other effect may arise from an informational asymmetry inherent in any decision-making process that involves the use of advice. Individuals are privy to their

own thoughts, but not to the thoughts underlying the advisor's opinion. A judge can access pieces of evidence supporting his/her own opinion more easily than ones supporting the advisor's view. If the weighting of opinions is a function of the accessible evidence, then, other things being equal, judges should be expected to discount advice.

A related hypothesis is that the weight of advice is a function of the judge's initial knowledge or competence. The more knowledgeable individuals are, the more evidence they retrieve from memory for their own opinion and, therefore, the higher weight they place on their own opinion.

Distance effects

How does discounting depend on the distance of the advice from one's own opinion? To develop the relevant hypotheses I used the aforementioned analogy between studies of advice use and of attitude change. In both situations individuals integrate their own prior opinion with that of another person. Research on the effects of influential messages on attitude change can inform us about how advice distance affects the way messages are weighted. Consider a practical advice-using situation in which your initial guess is that the distance between two places is roughly 10 miles. Then advisor A tells you she thinks the actual distance is 15 miles, while advisor B tells you he thinks the distance is 80 miles. The "near" advice might lead you to revise the initial estimate ("She says the place is somewhat further than I had initially thought"). The "far" advice, however, seems to call for a total reconsideration of the appropriate weighting strategy ("His opinion is too far from mine—either his or my estimate must be mistaken").

A basic tenet of social-cognitive psychology embedded in all consistency theories is that individuals seek to resolve discrepancies that exist among their beliefs. Theories of attitude change, such as dissonance (Aronson, Turner, & Carlsmith, 1963) and social judgment (Sherif & Hovland, 1961), predict that attitude change should decline with distance. Suppose attitude change is measured as a proportion—the amount of change is expressed as a fraction of the distance between the initial attitude and the message. Bochner and Insko (1966) presented a persuasive message advocating that people get N hours of sleep per night (where N ranged in various conditions from 8 to 0 h). The respondents' initial views (in an independent sample) averaged around 7 or 8 h per night. Then, as the advocated number of hours of sleep decreased—namely, the discrepancy increased—the magnitude of attitude change decreased. As the message becomes more extreme, people begin to generate counterarguments or disparage the source.

A related phenomenon was seen in studies of stereotype change (Kunda & Oleson, 1997), and conceptualized

in terms of assimilation and contrast processes (Sherif & Hovland, 1961). While a slightly deviant opinion could be assimilated and thus cause a shift in one's attitude, an extremely discrepant one has a proportionally reduced effect, since it falls outside the person's "latitude of acceptance" (Sherif & Hovland, 1961) and stands out in a stark contrast to one's initial opinion. The notion that social influence declines with distance has been incorporated in Davis et al.'s (1997) social judgment scheme. This model describes how the opinions of groups (e.g., committees, juries) are aggregated during discussion to establish the group's consensual judgment. An element of the model is the idea that a discrepant opinion's impact on group decision quickly declines as the discrepancy increases. In sum, the prediction based on attitude-change studies is that distant advice will be weighted less than near advice.

Using advice to improve accuracy

A major motivation for seeking advice is the expectation of improving judgment accuracy.¹ Numerous studies have indeed shown that combining multiple estimates tends to improve predictions (e.g., Armstrong, 2001; Ashton & Ashton, 1985; Libby & Blashfield, 1978; Sniezek & Buckley, 1995; Sniezek & Henry, 1989; Sorkin et al., 2001; Winkler & Poses, 1993; Yaniv, 1997; Yaniv & Hogarth, 1993; Zarnowitz, 1984).

A number of formal models provide a theoretical basis for understanding of when and how combining estimates improves accuracy. Accuracy is measured in terms of mean absolute error or judgment–criterion correlation. These include models based on the Condorcet jury theorem (majority rules/binary issues) and group signal-detection theory (Sorkin et al., 2001), models for combining subjective probabilities from multiple judges (Budescu & Rantilla, 2000; Wallsten, Budescu, Erev, & Diederich, 1997), and models for combining point forecasts (Clemen, 1989; Hogarth, 1978). In the case of quantitative judgments, a brief outline can show how and why improvement is to be expected from the use of advice. According to the Thurstonian view, a subjective forecast about an objective event is the sum of three components: the "truth," a constant bias, and random error. Statistical principles guarantee that forecasts formed by averaging several sources have lower variability than the individual opinions. The combined forecasts are expected to converge about the truth if the bias is zero or fairly small (e.g., Einhorn, Hogarth, & Klempner, 1977). In the present study we also investigate the effect of the respondents' weighting policies on the accuracy of their final judgments.

¹ Here the analogy between advice use and attitude change breaks down, since objective accuracy is not an issue in the study of attitudes.

Overview

In our studies we presented respondents with questions that had real consequences for them as decision makers, since they received a bonus for making accurate judgments. The respondents were given advice and the principal measure was the weight placed on that advice in their final decisions. The studies, which were conducted on a computer due to their interactive nature, shared the following general procedure. In the first phase, respondents were presented with questions and asked to state their estimates. In the second phase, they were presented with the same questions along with estimates made by various advisors (other students). The respondents were then asked to provide their estimates once again. They were free to use the advice as they wished. In Study 1 the advice was drawn at random from a pool of advice. In Studies 2 and 3 the advice was presented at one of three distance levels (near, intermediate, or far). Thus, the advice had to be “custom-made” on-line by the computer specifically for each respondent, depending on his or her initial opinions in the first phase.

Two important notes are in order. First, in all the studies we paid a bonus for each final estimate with a lower than average error, so it was in the respondents’ interest to consider carefully and make the best use of the advice in whatever manner they deemed appropriate. Second, a major advantage of the present experimental method (Studies 1–2) is the use of ecologically valid advice, that is, advice sampled from pools (distributions) of actual estimates made by other individuals. In the third study the advice was generated mechanically as a simple transformation of the respondents’ initial opinions. This method allowed us a certain control that could not be obtained in Study 2, at the expense of the ecological structure preserved in the first two studies. In sum, two alternative operational definitions of advice distance were tested. We compared the weighting policies, distance effects, and accuracy gains obtained using either the ecological or the mechanical advice.

Study 1: Weighting advice as a function of knowledge

The goal of the first study was to replicate and extend the discounting phenomenon and, in particular, to test whether advice discounting varies as a function of the judge’s knowledge. Such a finding would provide further support for our hypothesis. If discounting depends on evidence retrieval, then those who are more knowledgeable should place less weight on the advice than those who are less knowledgeable. Studies 2 and 3 further tested the interaction between advice distance and knowledge.

Method

The first study investigated how people use advice from a randomly drawn advisor in an ecological pool. The experimental procedure was conducted individually on personal computers. Fifteen questions about the dates of historical events (within the last 300 years) were presented sequentially on the computer display screen. As shown in Table 1, in the first phase respondents were shown one question at a time and asked to type in their best estimate for each one via the computer keyboard; in addition, they were asked to give lower and upper boundaries such that the true answer would be included between the limits with a probability of .95.

After the first phase was over, the respondents were told that there would be a second phase in which they would be presented with the same set of questions again. Now, however, each question would be presented along with two estimates: the respondent’s own initial estimate and that of an advisor. The respondents would then be asked to give a second, possibly revised, estimate for the question. No online feedback was given on the accuracy of their own or the advisors’ opinions (in particular, the correct answers were never shown). The respondents were told they would get a bonus at the end of the study, depending on their overall accuracy (see below).

The advisor’s estimate was randomly drawn by the computer from a pool of 50 estimates collected in an earlier study in which respondents were instructed merely to provide the best estimate for each question. The advisor varied from one question to the next, with labels such as A, D, and J used to indicate that each estimate came from a different individual. By sampling estimates from pools of data, adequate ecological validity could be maintained. The dispersion of the estimates and their errors corresponded to those that might have been encountered in reality by our respondents when seeking answers to such questions among their peers—undergraduate social science students.

The respondents ($N = 30$) were undergraduate students who participated either as part of their course requirements or for a flat fee of 12 Israeli shekels. They were all told that they would receive a bonus based on the accuracy of their estimates. In particular, they would

Table 1
Sample question and outline of the general procedure

Phase 1 (series of 15 questions): In what year was the Suez Canal first opened for use? Your best estimate ____ (low estimate ____ high estimate ____)
Phase 2 (same 15 questions repeated): In what year was the Suez Canal first opened for use? Your previous best estimate was <u>1905</u> The best estimate of advisor K was <u>1830</u> Your final best estimate ____

Table 2
Results from Study 1

Judge's knowledge	Weight of advice	Absolute error		% Improvement	Absolute error* weight+.17
		Before	After		
High	0.20	46.3	38.9	15	37.8
Low	0.33	66.0	50.7	21	48.1

*These are the mean absolute errors that would have been observed had respondents increased their actual weight of advice by 0.17 on every single trial.

receive 1 Israeli shekel (\$0.30 at the time of the study) as a bonus for each estimate that had a better than average accuracy score. Altogether they could collect up to 15 shekels in bonus payment. Thus it was in their interest to consider carefully and make the best use of the estimates given to them. The bonus was based on the final estimates (i.e., second phase).

Results

Advice weighting

The final estimate can be represented as a weighted combination of the two prior estimates—*initial* and *advice*—with the weights being proportional to the extent of the shift towards (away from) the advice. We define “weight of advice” = $|f - i|/|a - i|$, where i , f , and a stand for initial, final, and advice, respectively; the weight of advice is well-defined if the final estimate falls between the initial estimate and the advice, as it did in over 95% of the cases. The weight of advice, expressed as a proportion, reflects the weight that a respondent assigns the advice (and is inversely related to the extent to which the advice is discounted). Thus, the weight of advice takes a value of 0 if, in making the final estimate, the respondent adheres completely to his or her initial estimate (100% discounting of the advice); the weight of advice is 1.0 if the respondent shifts completely to the advice (0% discounting). Intermediate weights indicate that positive weights were assigned to both opinions (partial discounting).

Whereas a weight of 0.50 for advice implies equal weighting, the actual mean weight of advice (0.27) was significantly lower, $t_{29} = 6.35, p < .01$. Respondents placed a higher weight on their own opinion than on the advisor's opinion. This tendency was exhibited by most respondents: 28 of the 30 respondents had a mean weight of advice lower than 0.5. The respondents' means had an interquartile range from 0.19 to 0.47. Further analysis examined the distribution of all 450 individual trials (30 respondents \times 15 questions). After rounding to the nearest decimal, the weights of advice were classified into three groups: low (0–.3), medium (.4–.6), and high (.7–1.0). The percentages falling in these groups were 58%, 20%, and 22%, respectively. These results support the conclusion that individuals tend to discount advice.

Weighting as a function of personal knowledge

Next we analyzed the weight of advice as a function of the respondents' own knowledge, measured in terms of their prior performance. The respondents were divided into two groups (median split)—high knowledge and low knowledge—according to their accuracy (a function of average absolute error) in Phase 1 of the study, that is, depending on whether their average error fell below or above the median. As Table 2 shows, the high-knowledge group discounted the advice significantly more than the low-knowledge group. The respective mean weights of advice were 0.20 vs 0.33, $t_{28} = 2.65, p < .05$.

Improving accuracy

Exposure to the advice helped respondents improve their accuracy. The mean absolute error (in years) was reduced from 56.2 (for the initial estimate) to 44.8 for the combined estimate, $F(1, 28) = 14.02, p < .01$. Table 2 shows the accuracy gains for the two knowledge groups, 15% for the high-knowledge group (error reduced from 46.3 to 38.9 years), and 21% for the low-knowledge group (66.0–50.7). The low-knowledge group seemed to benefit more from the advice, but the interaction between knowledge group and type of error (initial vs final) was not significant, $F(1, 28) = 1.68$.

Discussion

A major conclusion regarding weighting policies in this study is that decision makers tend to discount advice. The respondents and the advisors were drawn from the same population, with similar background knowledge; on average, the respondents' accuracy was on a par with that of the advisors (mean absolute errors were 56.1 and 49.6, for respondents and advisors, respectively). Nevertheless, the respondents placed greater weight on their own judgments. They resolved the discrepancy between their own and the other opinion by adhering to their own opinion and making a token shift to the other opinion.²

² The results are similar to those obtained in Study 1 in Yaniv and Kleinberger (2000). The main difference between the two studies is that *no* feedback was given online in the present study (as noted in the method section), whereas in the previous study feedback—the correct answer—was given after each trial in the second phase, allowing respondents to track the accuracy of the advice and of their own estimates.

These results suggest two opposite perspectives on insight. On the one hand, the weights on advice were too low, suggesting that respondents' evaluations of their own knowledge were exaggerated overall. Indeed, people reveal poor insight in over-estimating the chances that their knowledge is correct (calibration curves reveal overconfidence, e.g., Lichtenstein & Fischhoff, 1977). On the other hand, respondents did not discount advice indiscriminately—those who knew less (in the first phase) placed higher weight on the advice than those who knew more. Such realism is also found in studies of probabilistic confidence judgment, where calibration curves are often found to be monotonically increasing, thereby indicating that easy items are assigned higher confidence levels than hard ones. Such findings indicate that self-assessment is not a unidimensional concept.

The advice-discounting hypothesis can explain both aspects of the present results. First, there is asymmetry in access to the evidence underlying each opinion, such that respondents are privy to their own thoughts but not those of the advisor, and therefore weight their own opinions more heavily. Second, those who know less presumably retrieve fewer pieces of evidence to support their estimate, so they tend to place higher weight on the advice (compared with those who know more).

To what extent was advice underweighted? To get a handle on this, we use as a first approximation the deviation of the average advice weight (0.27) from 0.50—a difference of 0.23. Another rough approximation of the amount of underweighting can be obtained empirically by calculating the “optimal” weight of advice on a trial-by-trial basis. The optimal weights were calculated assuming that the true answer for each question was known (hence a best weight of advice could be derived).³ The average optimal empirical weight of advice was 0.44, compared with the actual weight of 0.27, so the difference between them was 0.17.

To what extent might accuracy be improved if respondents increased the weight of advice? There are various ways to assess that potential improvement. The following calculation is given as an illustration. We calculated the final estimates that would have been obtained if the respondents had increased the weight assigned to the advice on each particular trial by 0.17 (the difference between the actual and the optimal weights). As Table 2 shows, the new final estimates were slightly more accurate than the actual final estimates (3–6% gain, not significant, $t_{29} = 1.66, p = .107$). Most of the gain in accuracy is already achieved by the respondents' actual final estimates. It seems that merely considering an additional opinion is the key to achieving greater accuracy, while its exact weighting is less critical. Studies of combining forecasts suggest that accuracy (or fit) is

highly robust to deviations of the weights from the optimum (Blattberg & Hoch, 1990).

Study 1 sets the stage for Studies 2 and 3, in which we systematically varied the distance of the advice from the decision maker's initial opinion. We asked whether and how respondents' weighting policy varies as a function of advice distance.

Study 2: Weighting ecological advice as a function of distance

We investigated how the distance of the advice from one's own opinion affects the weight it receives. The advice was at one of three distance categories: near, intermediate, or far. Each respondent experienced all three distance conditions, with one-third of the trials in each. The advisory estimates were designed online specifically for each respondent, depending on the estimates he or she gave in the first phase. For each question, the computer accessed a pool of estimates produced in previous studies and selected advice from it. This procedure guaranteed that estimates were selected from within the empirical distribution, and thus took into account the natural spread of the estimates. This design allowed us to test how people weight advice as a function of its distance from their initial opinions. In particular, we predicted that the greater the distance of the advice, the lower the weight it would be assigned. Moreover, we expected difference between high- and low-knowledge judges.

Method

Procedure

The procedure included two phases, as in Study 1. In the first phase the respondents ($N = 48$) were asked to produce estimates in answer to a list of questions. In the second phase they received the same list of questions along with advice and were instructed to provide their final estimates. There were a total of 24 trials with one-third of the questions in each of the three within-participant distance conditions: near, intermediate, and far. The three distance categories were presented in random order.

Selection of advice

For each question we had a pool of 120 estimates collected earlier in previous studies. For each respondent the computer generated advice for each question after phase 1 was over. The computer accessed the estimates for each question and sorted them in order of absolute distance (in years) from the respondent's point estimate (from the nearest to farthest). The advice to be offered to the respondent was then chosen according to its position relative to the initial estimate. In the near

³ The formula for deriving the weight of advice in estimating the true answer was similar to the one used in Study 1.

condition, the estimate in the 20th percentile was selected (i.e., the 24th nearest out of 120 estimates). In other words, 20% of the estimates were between the initial estimate and the advice. For the intermediate distance condition 55% of the estimates separated the initial estimate from the advice, and for the far distance condition the percentage was 90%. The mean absolute distance of the advice from the respondents' initial estimates was 24.1 in the near condition, and 50.1 and 93.8 in the intermediate and far conditions, respectively. The questions were randomly assigned to the different conditions.

Nothing was said to the respondents about how the estimates were selected from the pools. As in Study 1, we merely told respondents that the various pieces of advice were initial estimates generated by individuals who had participated in similar studies in the past. We also told them that at the end of the study they would be awarded a bonus for accuracy, 1 shekel (\$0.30 at the time of the study) for each estimate that had greater than average accuracy. Thus, they could earn up to 24 shekels in bonus payments altogether. Hence it was in their interest to consider their answers carefully and make the best use of the advice provided.

Results

As in the previous study, respondents were median-split into two knowledge groups according to their mean absolute error in the first phase of the study. The mean weights are shown in Table 3. An analysis of variance was performed on the weighting of advice with the decision maker's knowledge (high, low) and advice distance (near, intermediate, far) as factors. There were significant effects of knowledge, $F(1, 46) = 23.55, p < .001$ and distance, $F(2, 46) = 3.69, p < .05$, as well as an interaction, $F(2, 46) = 7.95, p < .01$. To understand the interaction, the simple effects were examined. The simple effect of knowledge was significant in the intermediate advice condition, $F(1, 46) = 26.8, p < .001$, and in the far advice condition, $F(1, 46) = 31.1, p < .001$, but not in the near condition $F(1, 46) = 1.65, p > .2$. In sum, the high-knowledge group generally placed less weight on the advice than did the low-knowledge group; moreover their weighting of the advice decreased with distance.

Table 3
Study 2: Weight of advice as a function of distance and the decision maker's knowledge

Decision maker's knowledge	Distance of advice		
	Near	Intermediate	Far
High	0.33	0.27	0.17
Low	0.44	0.53	0.49

Table 4
Study 2: Judgment errors before and after getting ecological advice

Decision maker's knowledge	Absolute error		% Improvement
	Before	After	
High	35	33	6
Low	64	47	27

The use of ecological advice improved accuracy by about 20%. The mean absolute errors before and after the advice was given were 50.1 and 40.2 years, respectively, $F(1, 46) = 48.9, p < .001$. As Table 4 shows, the accuracy gain was 6% for the high-knowledge group (error reduced from 35.4 to 33.2 years), and 27% for the low-knowledge group (from 63.5 to 46.7). This difference in accuracy gains led to a significant interaction between knowledge group and type of error (initial vs final), $F(1, 46) = 27.3, p < .001$.

Discussion

The high-knowledge respondents discounted the advice. Moreover, their weighting of the advice decreased systematically with distance. The low-knowledge group did not exhibit discounting nor did they display a clear pattern in weighting the advice, perhaps because they felt they could benefit even from distant advice (accuracy gains are shown in Table 4). We will return to this issue in the third study, in which the advice was generated differently.

In this study, advice was drawn from ecological samples of the estimates generated by other respondents in earlier studies. Advice distance was operationally defined relative to the natural distribution of the estimates given for each question, so that, for instance, far advice occupied the same relative position within the respective distributions. In our view, this design provides two important advantages. First, it helps make the advisory estimates seem realistic and believable, as having indeed been generated by other respondents. A second advantage is that the ecological design allows easier generalizations from experiment to reality. A disadvantage of the ecological design is that the absolute distances of the advice from the initial opinions could not be controlled. In particular, we did not control whether the advice was in the direction of the truth or leading away from the truth. In the next study we included this factor as well in the analysis.

Study 3: Weighting mechanical advice as a function of distance

In Study 3 the absolute distance of the advice was controlled. Advice was created mechanically by adding or subtracting a constant from the decision maker's

initial estimate. The use of advice that is a simple transformation of the initial estimates does not abide by the ecological constraints of the previous study, but it allowed us to test further our hypothesis regarding the influence of advice distance on weighting policies. We did this by separating the trials into two conditions, advice that was helpful (directed toward the truth) and not helpful (directed away from the truth). Thus we could analyze the effect of distance in either direction for the low- and high-knowledge groups.

Method

The procedure included the same two phases as in the previous study. In the first phase, the respondents ($N = 76$) were asked to produce estimates for 24 questions. In the second phase they received advice at various distances from their initial estimates and were asked to form their final estimates.

The procedure for generating the advice was as follows. Three sets of constants were created, based on the mean absolute distances in Study 2. The near advice was generated by either adding or subtracting one of the following constants from the initial estimates: 15, 18, or 20 years. The intermediate distance advice was generated at distances of 40, 43, or 45 years. The far advice was generated at distances of 70, 72, or 75 years. The use of three constants at each distance category was meant to obscure the underlying structure of the advice set (which indeed was not transparent to any of the respondents). Eight questions were randomly assigned to each of the three advice distance conditions (near, intermediate and far). The order of the various conditions was randomized for each respondent and the constants for creating the advice were sampled at random. The other aspects of the study were identical to those of the previous study, including the bonus for accuracy.

Results

The sample was median-split into two groups according to the respondents' degree of knowledge (a function of average absolute errors) in the first phase. The weight of the advice was calculated as in Study 1. Table 5 shows the mean weights as a function of the respondents' degree of knowledge and the advice distance. An analysis of

variance on the weights, with knowledge (high, low) and distance (near, intermediate, far) as factors, showed the following significance levels: distance, $F(2, 148) = 7.95, p < .005$, knowledge, $F(1, 74) = 3.89, p = .052$. Since knowledge was a significant factor in Studies 1–2, the one-tailed significance level $p < .05$ is warranted in this case. The interaction was not significant, $F < 1$. Specifically, the high-knowledge group discounted the advice more than the low-knowledge one. The weight of advice decreased as its distance from the initial opinion increased.

Next, the trials were separated into two conditions, according to the direction of the advice: helpful advice (pointing towards the truth) and unhelpful advice (pointing away from the truth). (There were half in each direction at each distance condition and for each respondent, by design.) The weights of the unhelpful advice were 0.29, 0.23, 0.18, for near, intermediate, and far, respectively, for the high-knowledge group, and 0.34, 0.28, 0.34, for the low-knowledge group. The respective weights of the helpful advice were 0.33, 0.33, 0.27, for the high-knowledge group, and 0.42, 0.40, 0.29, for the low-knowledge group. A three-way analysis of variance found significant effects of knowledge, $F(1, 74) = 4.18, p < .05$, direction, $F(1, 39) = 19.1, p < .05$, and distance, $F(2, 78) = 6.44, p < .05$. There were no significant two-way interactions, $F < 1$, but there was a significant triple interaction, $F(2, 148) = 3.36, p < .05$.

The effect of knowledge on the weight of advice was shown in previous analyses. The direction effect means that the helpful advice was weighted more than the unhelpful advice. Respondents presumably retrieved more support from their memory for the former type of advice. The declining pattern of weights in the high-knowledge condition was observed in both directions. The pattern of weights in the low-knowledge condition was not stable across directions. We will return to these results and the differences between Studies 2 and 3 in the final discussion.

In terms of accuracy, the mechanically generated advice was not as helpful to respondents as was the ecological advice in the first two studies. The mean absolute error barely changed as a result of receiving the advice (error reduced from 65.6 to 63.4), $F(1, 74) < 1$, yielding no significant accuracy gains, either overall or in one of the knowledge groups, as Table 6 shows. The results here greatly depart from those of Studies 1 and 2.

Table 5
Study 3: Weight of advice as a function of distance and the decision maker's knowledge

Decision maker's knowledge	Distance of advice		
	Near	Intermediate	Far
High	0.31	0.28	0.23
Low	0.38	0.34	0.30

Table 6
Study 3: Judgment errors before and after getting (mechanical) advice

Decision maker's knowledge	Absolute error		% Improvement
	Before	After	
High	52	49	6
Low	81	79	2

General discussion

We investigated two main aspects of advice use. The first involves the influence of advice on the decision maker's final judgment, and in particular the weight assigned to advice. The second involves the accuracy gains resulting from the weighting policy. We consider each of these aspects in turn.

Weighting advice

A coherent picture emerges from the advice weighting policies observed across the studies. First, the results of Study 1 show egocentric discounting of advice. Second, advice discounting was not indiscriminate; individuals had a veridical view of their knowledge, so that the less knowledgeable ones placed greater weight on the advice (Studies 1–3). Third, the weight of advice declined with the distance between the advice and their initial opinions (Studies 2–3); this distance effect was exhibited in the high-knowledge condition and to a lesser extent in the low-knowledge condition as well.

Advice discounting: A self/other effect

The asymmetric weighting of one's own and other opinions is attributed to the fundamental asymmetry in access to the underlying justifications for each opinion. Decision makers can assess what they know and the strength of their own opinions, but are far less able to assess what an advisor knows and the reasons underlying her/his opinions. Naturally, one's confidence about a given opinion (or hypothesis) is related to the amount of evidence that one could readily recruit to support it. Other things being equal, decision makers are likely to feel more confident about their own opinion than about the other opinion, hence their own estimate would receive greater weight than the advice. Earlier findings suggest that respondents weight each opinion according to the expertise ascribed to its source (Birnbaum & Stegner, 1979; Birnbaum & Mellers, 1983). The self/other asymmetry presumably enhances the expertise ascribed to the self. This line of reasoning about information asymmetry is also reminiscent of the principal-agent problem in organizations (Eisenhardt, 1989).

There is also other evidence for advice discounting. Harvey and Fischer (1997), using a cue-learning task, had respondents make initial estimates and then final estimates on the basis of a recommendation from an advisor. They found a shift in judgment of about 20–30% towards the advice—a result consistent with what we observed. Using a time-series forecasting task, Lim and O'Connor (1995) had respondents integrate a statistical forecast into their initial judgment-based forecast. These respondents assigned about double the weight to their own initial forecast than to the statistical

forecast. Sorkin et al. (2001) also report higher weights placed on one's own opinion in a group signal-detection task. On each trial, one member was randomly selected and told that she was to give the group's answer, on the basis of the other members' responses. A participant's weight was consistently higher when she was the designated responder.

There is evidence that such discounting also occurs in professional settings. In his literature review on the impact of genetic counseling, Kessler (1989) concludes that genetic counseling does not produce dramatic changes in counselees' reproductive decisions. The best predictor of the post-counseling reproductive decision is the counselee's pre-counseling intentions. Advice discounting may also be related to the public's perception of risks (such as environmental and health-related risks). A recurring finding is that experts and the public differ in their perception of such risks, thus hindering the implementation of public policy (Flynn, Slovic, & Mertz, 1993). Experts' risk communication can be viewed as advice to individuals in their daily decisions regarding the safety measures they need to take against various types of risks (e.g., radiation from mobile phones, using a mobile phone while driving). The observed skepticism towards expertise can be viewed as a form of discounting of the experts' advice. Finally, the phenomenon that individuals stick closely to their initial opinions is also consistent with the findings of perseverance and resistance to change known from classical research on attitudes (e.g., Sherman & Cohen, 2002).

Alternative accounts

Motivational effects

The explanation of the self/other effect in terms of differential information access seems preferable to alternative explanations that posit either a self-serving bias (e.g., an optimistic bias) or commitment to one's past decisions as the root of discounting others' views. To be sure, self-serving biases pervade interpersonal comparisons, in that, for example, people believe that they have lower chances of experiencing negative life events, such as car accidents and strokes, than others do or that they rank higher than others on various abilities and attributes, such as driving ability and social skills (e.g., Brown, 1986). But a bias of this sort does not readily explain respondents' weighting policies for advice, especially the sensitivity of those policies to the respondents' own knowledge (Studies 1–3) and their sensitivity to the quality of the advice (Yaniv & Kleinberger, 2000).

Commitment to one's past decisions is a powerful motive in decision making, yet it cannot readily explain the findings either. The antecedents of commitment—high costs for being inconsistent, the need to justify decisions to others, having to admit past mistakes, and having to save

face with respect to ego-involving issues—were largely absent in the present studies. Our respondents made their judgments in a private setting (by entering responses into a computer file), received incentives for accuracy, and were not asked to justify their estimates.

A cognitive explanation based on informational asymmetry and the assessment of available evidence is more parsimonious and hence superior to those based on a self-serving bias or commitment because it can readily account for the finding that respondents' weights on advice are sensitive to the quality of the advice (Yaniv & Kleinberger, 2000) as well as their own knowledge (e.g., Study 1), without making unnecessary assumptions.

Information integration

Our account of the present results on weighting advice is linked to theories in the tradition of information integration. Such theories posit simple cognitive processes to explain the updating of impressions and beliefs. Anderson (1968) attributes the primacy effect in impression formation to attention decrement over successive serial positions as the weights given to later cues in a sequence decrease. Expanding on such ideas, Hogarth and Einhorn (1992) introduced a formal model of how people update their beliefs on the basis of sequential information (e.g., pieces of evidence in a trial or a list of personality traits). A central characteristic of the updating process, according to Hogarth and Einhorn, is the response mode, namely whether updating is made globally at the end of the sequence or step by step, after each item is presented.

According to Hogarth and Einhorn, the end-of-sequence mode is conducive to primacy effects, and the step-by-step mode, to a recency effect. Our respondents' behavior shows a primacy effect, as they preferred their own opinion to the advice (e.g., Study 1). In this respect our findings are in agreement with the prediction for the end-of-sequence mode, based on the belief-updating model. But our decision-advice-revise procedure does not fall squarely into either of the response mode categories—"end of sequence" or "step by step"—since respondents had in fact generated one of the two estimates themselves in an earlier phase. This differs from information integration studies, where the sequences of items are fully controlled by the experimenter. Moreover, the sequential nature of the belief-updating model makes the order of presentation a key factor. Our procedure highlights the judge's own opinion, hence order, being just one factor among others, may not necessarily be as important a factor as self/other asymmetry.

The present studies like the information-integration approach, focus on respondents' weighting policies. The present studies highlight additional key features, including the use of realistic (rather than fictional) information, thereby enabling respondents to rely on

pre-experimental knowledge. In sum, I suggest that our decision-advice-revise procedure adds another aspect to information integration, one which has not been explored so far, and is potentially fruitful.

The effect of advice distance on the revision of opinion

We hypothesized that the weight of advice would decline as its distance from the respondent's initial opinion grew larger. It appears that knowledge modulates the distance effect. The decline of weight with distance was shown consistently for the high-knowledge respondents (in Studies 2–3), but less regularly for the low-knowledge respondents (in Study 3, but not in Study 2). While we did not predict a difference between high- and low-knowledge respondents, we can make sense of these findings.

The more knowledgeable individuals presumably have a narrower latitude of acceptance than the less knowledgeable individuals. Therefore the two groups differ in their attributions. The more knowledgeable judges, according to this hypothesis, are more likely to attribute the discrepancy between their own and another person's opinion to the other person's fault or error rather than their own. In particular, upon encountering a different opinion the two groups proceed with different inferences—the more knowledgeable respondents with "I guess *the advisor* is wrong" and the less knowledgeable ones with "I guess *I am* wrong." Such attributions might evolve from the respondents' different experiences. The initial views of the knowledgeable judges are often in the neighborhood of the best solution, hence they tend to assume that near advice is of good quality while far advice is of lower quality. The less knowledgeable judges might be less inclined to use distance as a predictor of advice quality, since their own hunches are less accurate. This might explain why the distance effect was less pronounced for the low-knowledge respondents.

The present findings on the distance effect are consistent with earlier work on attitude change which suggests that the influence of a message (measured as a proportional change) tends to decrease as a function of its discrepancy from the recipient's initial attitude (Bochner & Insko, 1966).⁴ In a more recent work on stereotype change, Kunda and Oleson (1997) tested the influence of a single counter-stereotypic example on existing personal stereotypes. For instance, given the stereotype that public relations (PR) people are extroverts, Kunda and Oleson presented to respondents

⁴ In fact, this could lead to a phenomenon called the boomerang effect. When a message is highly discrepant, judges shift their attitude less toward it than they would have done had the message been less discrepant.

either an extremely deviant example (i.e., an extremely introverted PR person) or a moderately deviant example (a slightly introverted PR person). The extreme example had less influence on stereotype change, in accord with predictions of assimilation/contrast theories (Sherif & Hovland, 1961). Specifically, a slightly deviant example is easily assimilated into the stereotype and hence can change it, whereas an extremely discrepant one is in great contrast to the stereotype and so is likely to be discounted. Recent work on anchoring has also shown that extreme anchors have proportionally less effect on judgment than moderate ones (Marti & Wissler, 2000; Wegener, Petty, Detweiler-Bedell, & Jarvis, 2001). According to these authors, judges tend to discredit or argue against extreme anchors, thereby making them less influential.

In a different vein, early studies on information integration (Anderson & Jacobson, 1965) and additive models in judgment (Slovic, 1966) suggest that judges discount inconsistent cues. Moreover studies of the process of combining opinions show that judges give greater weight to consensus opinions while discounting outlier opinions (Yaniv, 1997). Finally, studies of group decision-making suggest that a discrepant opinion's impact on the group's final decision declines as the discrepancy increases (Davis et al., 1997). In the foregoing studies an opinion (or cue) is discounted due to its distance from the *consensus*. In contrast, in the studies reviewed above an opinion is discounted due to its distance from the judge's *initial opinion*. The common thread between the two phenomena is that inconsistent information is discounted.

The benefit of advice

By consulting one advisory opinion—randomly sampled from an ecological pool of estimates—individuals in Study 1 improved their estimation accuracy by about 20%. There is a straightforward important consequence of such findings which often escapes people's attention. People do not always realize that in order to be helpful, the other opinion need *not* come from a smarter or more knowledgeable individual than the decision makers themselves. To reap the accuracy gains from aggregation, the additional opinions only need come from independent advisors (though small deviations from perfect independence still permit appreciable gains; e.g., Johnson, Budescu, & Wallsten, 2001).

That combining opinions improves accuracy is one of the most robust findings in the judgment literature. The explanation for the observed accuracy gains in the present studies was outlined briefly earlier—it relies (as all formal models do) on the central limit theorem in statistics as well as certain empirical facts about the task, such as the bias and inter-judge correlations (e.g., Wallsten et al., 1997; Johnson et al., 2001). Indeed, the

results of Study 2 also show accuracy gains. In Study 3, the advice was generated mechanically, by arbitrarily adding or subtracting a constant from the original opinion in the first phase. Since the advice was highly correlated with the initial opinion, we did not expect accuracy gains in that study.

Receiving and using other types of advice

The present research involved *quantitative* advice about *factual* matters (dates of events). Future investigations could and should be extended to include other types of advice. One might distinguish between *qualitative* (verbal) advice and quantitative advice. In particular, verbal advice does not lend itself to the same sort of weighting evaluated in the present studies.

In addition, one could distinguish between opinions about matters of *fact* (estimates or forecasts) and about matters of *taste* (evaluations or attitudes). The benefit accrued from combining opinions about matters of fact is both demonstrable and understood theoretically. In contrast, simple aggregation of tastes for the purpose of individual decision making—such as opinions about a movie that one has not seen or a restaurant that one has yet to try out—raises conceptual difficulties. People are entitled to their different tastes and it is less clear how individuals might combine their own preferences with those of a friend, colleague, or professional advisor. Thus a theory about combining opinions in matters of taste is in order. A related question is whether consulting others' opinions about matters of taste helps improve decision quality (assuming an acceptable definition of quality).

The present perspective suggests ways of thinking about how these other types of advice might be integrated. I suggest that qualitative advice, such as opinions about taste, helps decision makers overcome certain common weaknesses in reasoning. The relevant weaknesses include decision makers' failure to generate enough alternatives for choice and their tendency to try to confirm rather than disconfirm their prior views. For example, Svenson's (1996) differentiation-consolidation theory claims, in the tradition of dissonance theories, that self-confirmation is an ongoing, continuous process through which individuals construct justifications for their decisions.

I suggest that receiving advice (of any type) serves an adaptive function since it helps individuals overcome self-confirmation tendencies. Advisors can expose decision makers to unattended alternatives and unintended consequences, thereby challenging them to rethink their prior opinions and weigh the new and different opinions using some sort of an internal negotiation process that eventually yields a compromise between the two opinions. I do not claim that advisors are free of reasoning biases, but rather that, being independent, they

effectively challenge the decision makers with ideas that they might not gather *on their own* otherwise. Related suggestions appear, for instance, in Jonas and Frey's (2003) findings that advisors conduct a balanced information search and, under certain conditions, transmit both confirming and disconfirming information to personal decision makers. Last but not least, it appears that a most promising avenue for studying further the impact and benefit of advice about matters of taste involves the role of the "personal match" between the givers and receivers of advice. Presumably the greater the perceived similarity in characteristics (e.g., traits, background, and education), the greater the impact and benefit of receiving the advice.

In sum, researchers of individual decision-making have traditionally developed and investigated various decision-support systems that might help individuals improve their decisions (decision trees, formal models, computer models, etc.). I suggest that the social-cognitive function of seeking advice as a "corrective procedure" or support system for the *individual decision maker* has not been explored sufficiently. It is not surprising that advice-seeking pervades daily decisions, ranging from the choice of a movie to a decision about the promotion of an employee. What is surprising is that so little attention has been paid in decision research to a process so fundamental in real life. It is imperative for future research to consider the procedures by which various type of advice (e.g., qualitative verbal advice, opinions about matters of taste) are elicited and used best.

References

- Ajzen, I. (2001). Nature and operation of attitudes. *Annual Review of Psychology*, *52*, 27–58.
- Anderson, N. H. (1968). Application of a linear-serial model to a personality-impression task using serial presentation. *Journal of Personality and Social Psychology*, *10*, 354–362.
- Anderson, N. H., & Jacobson, A. (1965). Effect of stimulus inconsistency and discounting instructions in personality impression formation. *Journal of Personality and Social Psychology*, *2*, 531–539.
- Armstrong, J. S. (2001). *Principles of forecasting: A handbook for researchers and practitioners*. Dordrecht, Netherlands: Kluwer.
- Aronson, E., Turner, J., & Carlsmith, M. (1963). Communicator credibility and communicator discrepancy as determinants of opinion change. *Journal of Abnormal and Social Psychology*, *67*, 31–36.
- Ashton, A. H., & Ashton, R. H. (1985). Aggregating subjective forecasts: Some empirical results. *Management Science*, *31*, 1499–1508.
- Birnbaum, M. H., & Mellers, B. A. (1983). Bayesian inference: Combining base rates with opinions of sources who vary in credibility. *Journal of Personality and Social Psychology*, *45*, 792–804.
- Birnbaum, M. H., & Stegner, S. E. (1979). Source credibility in social judgment: Bias, expertise, and the judge's point of view. *Journal of Personality and Social Psychology*, *37*, 48–74.
- Blattberg, R. C., & Hoch, S. J. (1990). Database models and managerial intuition: 50% model + 50% manager. *Management Science*, *36*, 887–899.
- Bochner, S., & Insko, C. A. (1966). Communicator discrepancy, source credibility, and opinion change. *Journal of Personality and Social Psychology*, *4*, 614–621.
- Brown, J. D. (1986). Evaluations of self and others: Self-enhancement biases in social judgments. *Social Cognition*, *4*, 353–376.
- Budescu, D. V., & Rantilla, A. K. (2000). Confidence in aggregation of expert opinions. *Acta Psychologica*, *104*, 371–398.
- Clemen, R. T. (1989). Combining forecasts: A review and annotated bibliography. *International Journal of Forecasting*, *5*, 559–583.
- Davis, J. H., Zarnoth, P., Hulbert, L., Chen, X.-p., Parks, C., & Nam, K. (1997). The committee charge, framing interpersonal agreement, and consensus models of group quantitative judgment. *Organizational Behavior and Human Decision Processes*, *72*, 137–157.
- Einhorn, H. J., Hogarth, R. M., & Klempner, E. (1977). Quality of group judgment. *Psychological Bulletin*, *84*, 158–172.
- Eisenhardt, K. (1989). Agency theory: An assessment and review. *Academy of Management Review*, *14*, 57–74.
- Flynn, J., Slovic, P., & Mertz, C. K. (1993). Decidedly different: Expert and public views of risks from a radioactive waste repository. *Risk Analysis*, *13*, 643–648.
- Harvey, N., & Fischer, I. (1997). Taking advice: Accepting help, improving judgment and sharing responsibility. *Organizational Behavior and Human Decision Processes*, *70*, 117–133.
- Hogarth, R. M. (1978). A note on aggregating opinions. *Organizational Behavior and Human Performance*, *21*, 40–46.
- Hogarth, R. M., & Einhorn, H. J. (1992). Order effects in belief updating: The belief-adjustment model. *Cognitive Psychology*, *24*, 1–55.
- Johnson, T. R., Budescu, D. V., & Wallsten, T. S. (2001). Averaging probability judgments: Monte Carlo analyses of asymptotic diagnostic value. *Journal of Behavioral Decision Making*, *14*, 123–140.
- Jonas, E., & Frey, D. (2003). Information search and presentation in advisor-client interactions. *Organizational Behavior and Human Decision Processes*, *91*, 154–168.
- Jungermann, H. (1997). When you can't do it right: Ethical dilemmas of informing people about risks. *Risk Decision and Policy*, *2*, 131–145.
- Kennedy, J., Kleinmuntz, D. N., & Peecher, M. E. (1997). Determinants of the justifiability of performance in ill-structured audit tasks. *Journal of Accounting Research*, *35*, 105–123.
- Kessler, S. (1989). Psychological aspects of genetic counseling: A critical review of the literature dealing with education and reproduction. *American Journal of Medical Genetics*, *34*, 340–353.
- Kunda, Z., & Oleson, K. C. (1997). When exceptions prove the rule: How extremity of deviance determines the impact of deviant examples on stereotypes. *Journal of Personality and Social Psychology*, *72*, 965–979.
- Libby, R., & Blashfield, R. K. (1978). Performance of a composite as a function of the number of judges. *Organizational Behavior and Human Performance*, *21*, 121–129.
- Lichtenstein, S., & Fischhoff, B. (1977). Do those who know more also know more about how much they know? *Organizational Behavior and Human Performance*, *20*, 159–183.
- Lim, J. S., & O'Connor, M. (1995). Judgmental adjustment of initial forecasts: Its effectiveness and biases. *Journal of Behavioral Decision Making*, *8*, 149–168.
- Marti, M. W., & Wissler, R. L. (2000). Be careful what you ask for: The effect of anchors on personal injury damages awards. *Journal of Experimental Psychology: Applied*, *6*, 91–103.
- Sherif, M., & Hovland, C. I. (1961). *Social judgment: Assimilation and contrast effects in communication and attitude change*. New Haven, Ct: Yale University Press.

- Sherman, D. K., & Cohen, G. L. (2002). Accepting threatening information: Self-affirmation and the reduction of defensive biases. *Current Directions in Psychological Science*, *11*, 119–122.
- Slovic, P. (1966). Cue-consistency and cue-utilization in judgment. *The American Journal of Psychology*, *79*, 427–434.
- Sniezek, J. A., & Buckley, T. (1995). Cueing and cognitive conflict in judge-advisor decision making. *Organizational Behavior and Human Decision Processes*, *62*, 159–174.
- Sniezek, J. A., & Henry, R. A. (1989). Accuracy and confidence in group judgment. *Organizational Behavior and Human Decision Processes*, *43*, 1–28.
- Sorkin, R. D., Hayes, C. J., & West, R. (2001). Signal detection analysis of group decision making. *Psychological Review*, *108*, 183–203.
- Svenson, O. (1996). Decision making and the search for fundamental psychological regularities: What can be learned from a process perspective? *Organizational Behavior and Human Decision Processes*, *65*, 252–267.
- Wallsten, T. S., Budescu, D. V., Erev, I., & Diederich, A. (1997). Evaluating and combining subjective probability estimates. *Journal of Behavioral Decision Making*, *10*, 243–268.
- Wegener, D. T., Petty, R. E., Detweiler-Bedell, B. T., & Jarvis, W. B. G. (2001). Implications of attitude change theories for numerical anchoring: Anchor plausibility and the limits of anchor effectiveness. *Journal of Experimental Social Psychology*, *37*, 62–69.
- Winkler, R. L., & Poses, R. M. (1993). Evaluating and combining physicians' probabilities of survival in an intensive care unit. *Management Science*, *39*, 1526–1543.
- Yaniv, I. (1997). Weighting and trimming: Heuristics for aggregating judgments under uncertainty. *Organizational Behavior and Human Decision Processes*, *69*, 237–249.
- Yaniv, I., & Hogarth, R. M. (1993). Judgmental versus statistical prediction: Information asymmetry and combination rules. *Psychological Science*, *4*, 58–62.
- Yaniv, I., & Kleinberger, E. (2000). Advice taking in decision making: Egocentric discounting and reputation formation. *Organizational Behavior and Human Decision Processes*, *83*, 260–281.
- Yates, J. F., Price, P. C., Lee, J., & Ramirez, J. (1996). Good probabilistic forecasters: The "consumer's" perspective. *International Journal of Forecasting*, *12*, 41–56.
- Zarnoth, P., & Sniezek, J. A. (1997). The social influence of confidence in group decision making. *Journal of Experimental Social Psychology*, *33*, 345–366.
- Zarnowitz, V. (1984). The accuracy of individual and group forecasts from business and outlook surveys. *Journal of Forecasting*, *3*, 11–26.
- Zimbardo, P. G., & Leippe, M. R. (1991). *The psychology of attitude change and social influence*. Philadelphia: Temple University.

Received 27 June 2002