

האוניברסיטה העברית בירושלים

THE HEBREW UNIVERSITY OF JERUSALEM

SPURIOUS CONSENSUS AND OPINION REVISION: WHY MIGHT PEOPLE BE MORE CONFIDENT IN THEIR LESS ACCURATE JUDGMENTS?

by

**ILAN YANIV, SHOHAM CHOSHEN-HILLEL
and MAXIM MILYAVSKY**

Discussion Paper # 492

November 2008

מרכז לחקר הרציונליות
**CENTER FOR THE STUDY
OF RATIONALITY**

Feldman Building, Givat-Ram, 91904 Jerusalem, Israel
PHONE: [972]-2-6584135 FAX: [972]-2-6513681
E-MAIL: ratio@math.huji.ac.il
URL: <http://www.ratio.huji.ac.il/>

To appear in: *Journal of Experimental Psychology: Learning, Memory, and Cognition*

**Spurious Consensus and Opinion Revision:
Why Might People Be More Confident in Their Less Accurate Judgments?**

Ilan Yaniv, Shoham Choshen-Hillel, Maxim Milyavsky

The Hebrew University of Jerusalem

Abstract

In the interest of improving their decision-making, individuals revise their opinions on the basis of samples of opinions obtained from others. However, such a revision process may lead decision-makers to experience greater confidence in their less accurate judgments. We theorize that people tend to underestimate the informative value of independently drawn opinions, if these appear to conflict with one another, yet place some confidence even in the “spurious consensus” which may arise when opinions are sampled interdependently. The experimental task involved people’s revision of their opinions (caloric estimates of foods) on the basis of advice. The method of sampling the advisory opinions (independent or interdependent) was the main factor. The results reveal a dissociation between confidence and accuracy. A theoretical underlying mechanism is suggested whereby people attend to consensus (consistency) cues at the expense of information on interdependence. Implications for belief-updating and for individual and group decisions are discussed.

This research was supported by Grant No. 344/05 from the Israel Science Foundation to the first author. Address correspondence to Ilan Yaniv, Department of Psychology & Center for the Study of Rationality, Hebrew University of Jerusalem, Israel, Email: ilan.yaniv@huji.ac.il.

“No one wants advice – only corroboration.” – John Steinbeck

“Conflict is the gadfly of thought. It stirs us to observation and memory... It shocks us out of sheep-like passivity, and sets us at noting and contriving.” – John Dewey

In the interest of improving their decision-making, individuals often rely on the naïve (word-of-mouth) advice of their friends or neighbors (Harvey & Fischer, 1997; Yaniv, 2004a). In particular, they revise their beliefs on the basis of samples of opinions with the goal of improving their judgment. At the focus of this research is the process by which advisory opinions are *sampled* – an issue of great importance which bears on the perceived consistency of the opinions, the interdependence among them and, as we shall see, the performance of the users of the advice.

We suggest that decision-makers are prone to feel more confidence in their less accurate (advice-based) judgments, to the extent that they pay attention to the consistency of the advice and they fail to appreciate the consequences of interdependence among the opinions (a condition called “spurious consensus”). The dissociation between confidence and accuracy arises because decision-makers feel confident in their judgments when these are based on consensual opinions (albeit interdependent ones), yet underestimate the informative value of independent opinions that are conflictual.

In this paper, we review cognitive and social factors that might confound people’s inferences in everyday settings, involving the influence of the sampling method on the internal consistency and interdependence among the opinions. We then report a study investigating the conditions for the dissociation. In conclusion, we consider the adverse effects of incorrect inferences on the quality of opinion-revision and advice-based decision-making. Potential links to group decision-making are also pursued.

Consensus, Conflict, and Interdependence

A salient characteristic of a given set of advisory opinions is their perceived internal consistency, that is, the amount of agreement among them, and their agreement with the decision-maker’s prior opinion. Decision-makers’ ease of processing and confidence in their judgments increases as a function of the consistency of the information profiles with which they are presented (Budescu & Yu, 2006; Weber, Bockenholt, Hilton, & Wallace, 2000). Decision-makers find agreement with advisors rewarding because it confirms their initial choice, thus reducing their need to put more cognitive effort into the task (Savadori, Van Swol, & Snizek, 2003). Furthermore, decision-makers maintain consistency by giving greater weight to consensus opinions (Harries, Yaniv, & Harvey, 2004; Yaniv, 1997) and by assigning greater

weight to opinions that are consistent with their own prior opinion (Yaniv, 2004a, 2004b; Yaniv & Milyavsky, 2007).

The idea that consensus of opinions enhances decision-makers' confidence in their decisions seems like an uncontested descriptive hypothesis. However, a normative analysis qualifies the link between consistency and confidence in important, insightful ways. Basic statistical principles suggest that the benefit of having additional opinions is a function of the *level of dependence* among them (Budescu & Yu, 2007; Johnson, Budescu, & Wallsten, 2001). Intuitively, the level of interdependence between advisors A and B is the extent to which knowing A's opinion helps one guess B's opinion. Highly interdependent opinions are redundant and thus do not give the decision-maker new information. Other things being equal, one's confidence in one's judgment should be greater if it is based on independent rather than interdependent opinions.

Formal analysis demonstrates the superiority of independent opinions over dependent ones (Hogarth, 1978). Appreciable gains from combining the opinions of several judges should occur if their judgment errors are weakly correlated (e.g., $r = .1$), whereas only minute marginal gains should occur when judges' errors are highly correlated (e.g., $r = .9$). When interdependence is high, most of the expected gains from additional opinions are obtained with very few opinions (Yaniv & Milyavsky, 2007).

Interdependence among sources of advice is not uncommon in social settings. First, advisors may influence each other or, if they are members of a hierarchical organization, then they may be subject to the same influences. Either way the opinions become more interdependent. Second, advisors who rely on similar data and employ similar methods of interpreting the data are likely to reach correlated opinions, even if there is no social contact among them (e.g., unrelated financial analysts who subscribe to the same databases and software may reach similar conclusions). Finally, decision-makers may preferentially solicit, intentionally or unconsciously, the opinions of those who subscribe to their own position (Lord, Ross, & Lepper, 1979). Thus their biased sampling essentially inflates the interdependence among their sources.

These scenarios illustrate how the nature of the environment and the actions of the advisors and the decision-maker lead to interdependent sampling and hence spurious agreement among the sampled opinions. The foregoing analysis suggests that a valid consensus (i.e., a consensus that arises from independent sources) warrants increased confidence, while a spurious consensus (i.e., a consensus among opinions produced by interdependent sources) does not. Thus, individuals should increase their confidence on the basis of consensus only when the advisors are independent, *not* when they are highly interdependent. Do people recognize the

detrimental effects of great interdependence among sources? Do they sufficiently discount their confidence when told that the sources are redundant?

Previous studies suggest that people have only a fragmented understanding of the concept of independence. Notably, Kahneman and Tversky (1973) conjectured that “high intercorrelations among inputs increase confidence and decrease validity.” They also coined the term *illusion of validity*, suggesting that people do not recognize the role of independence. Kroll, Levy, and Rapoport (1988) investigated participants’ behavior in a task involving the prediction of stock prices based on forecasts made by analysts. The participants were influenced by the accuracy of the analysts’ forecasts, but not by the interdependence among them. Either the participants did not detect this interdependence, or else they used it incorrectly.

Maines (1996) presented participants with information on past intercorrelations among forecasters. She also found that individuals’ combined forecasts were sensitive to the forecasters’ relative accuracy, but not to the dependence among them. Soll (1999) asked his participants to indicate which of two hypothetical sources they would consult, given the one opinion they already had on hand. The participants also revealed a limited understanding of the relevance of independence, which could be traced to their implicit theories about the sources of judgment errors. Finally, Budescu and Yu’s (2007) participants were also singularly insensitive to inter-cue correlations, although they showed sensitivity to other task parameters.

Confidence-Accuracy Dissociation

The ideas and findings reviewed so far suggest dissociation between confidence and accuracy in the opinion-revision process. Settings that induce greater confidence may not necessarily lead to better performance; likewise, settings that lead to better performance may not induce greater confidence. If people are highly sensitive to consensus cues and less so to interdependence cues, then their confidence may increase after revision, even if their accuracy does not. Such dissociation is likely in settings involving a “spurious consensus” – that is, a set of agreeing opinions produced by inter-dependent sources. Similarly, if people underestimate the informational value of independent opinions (which, under uncertainty, tend to disagree with each other), then they may not feel any more confident after integrating independent opinions, even though their accuracy is likely to improve in the process (e.g., Yaniv, 2004b). Our research investigates this possibility.

A novel aspect of the present research is our manipulation of the method by which the opinions presented to the participants are sampled. We investigated how individuals use opinions generated by different sampling methods. With the first method, called *independent* sampling of advice, sets of opinions are randomly drawn from ecological pools of opinions.

With the second method, called *opinion-dependent* sampling of advice, sets of opinions are selectively drawn from those that agree with the decision-maker's initial opinion. The latter method generates sets of opinions that are more internally consistent than the former method, but also more interdependent.

Experiment

The experimental task involved estimating the number of calories in measured quantities of different foods (e.g., a cup of yogurt, a bowl of cooked rice). In the first phase of the procedure (shown in Table 1), participants were asked to generate a calorie estimate for each food and then indicate their confidence in it. In the second phase, they were provided with the opinions of three advisors, and were given the opportunity to revise their initial estimates. They were told that they would receive a bonus for making accurate judgments, so it was in their interest to make the best use of the advice, in whatever way they thought was appropriate. They were also asked to indicate their confidence in their final (revised) estimates and to bet on their accuracy.

Unlike previous research, we used real rather than hypothetical questions to which answers could be right or wrong and we elicited consequential judgments carrying rewards. This allowed us to measure confidence in relation to accuracy. Also, the use of the sampling method as an experimental factor allowed us to create independent and opinion-dependent samples of advice that were all drawn from ecologically valid pools of opinions. Participants were specifically told how the opinions were sampled on each trial. We analyzed participants' accuracy gains, confidence ratings in their initial and final estimates, and willingness to make bets on their answers.

Method

Procedure. The experimental procedure was conducted individually on personal computers, and included two phases. In the first phase, 28 questions on the caloric value of various foods were presented. Participants ($n=44$) were shown one question at a time and were asked to type in their best estimate for each. They were told that they would get a token compensation (equivalent to about \$0.03) for each estimate that was "close to the truth" (i.e., an estimate in the neighborhood of the correct answer, within $\pm 12\%$). Participants were also asked to rate their confidence in their estimate on a scale anchored at 0% (not confident at all) and 100% (completely confident), with a tick point at every 10%.

Table 1

*Outline of the General Procedure*Phase 1 (series of questions):

What is the calorie value of one serving of 3% yogurt?

Your best estimate _____

Phase 2 (same questions repeated):

What is the calorie value of one serving of 3% yogurt?

Your previous best estimate was 350

The best estimate of advisor #12 was 415

The best estimate of advisor #19 was 330

The best estimate of advisor #82 was 335

Your final best estimate _____

In the second phase of the procedure the participants were presented with the same set of 28 questions. Now, however, each question was presented along with the participant's own estimate (from the initial phase) and three advisory estimates. On each trial a header appeared on the screen, indicating whether the advice was randomly sampled or selected from the closest estimates. Specifically, on half the trials (independent condition) the header stated that "these estimates were *randomly drawn* from a pool of 100 estimates made by participants in a previous study," whereas on the remaining trials (opinion-dependent condition) the header stated that "these estimates were selected from those *closest to your own initial opinion* in a pool of 100 estimates made by participants in a previous study."

The participants were asked to give a final, possibly revised, estimate for the answer. They were again asked to rate how confident they were in their answer, and then, in addition, were asked to decide whether they would bet on it for a larger, 2-shekel bonus (about \$0.60) that they would earn if the answer fell in the neighborhood ($\pm 12\%$) of the correct answer. They were told they had to bet online on a total of 14 out of 28 answers. To assist the participants in tracking their bets, two counters appeared on the screen. One showed the number of the remaining bets (out of 14), and the other, the number of the remaining questions (out of 28).

Materials. In the independent condition, the participants were actually given three advisory estimates which were randomly drawn by the computer from a pool of 100 estimates. These estimates had been collected in an earlier study in which participants provided their best calorie estimate for each of the foods on the list. For each question, new advisors were sampled

at random. Thus the advisors varied from one question to the next, with labels such as #19, #80, and #2 indicating that the estimates came from different individuals on each trial.

In the opinion-dependent sampling condition, the three advisory opinions presented on each trial were constructed online for each participant, depending on his or her initial estimate in the first phase. The computer rank-ordered all the estimates for a given question, according to their distance from the participant's initial estimate, from the 1st (nearest) to 100th (farthest). The computer then selected and presented the three opinions in locations 1, 7, and 15.

The level of consensus among the four opinions (self plus three advisory opinions) presented on each trial of the second phase was assessed in terms of the "coefficient of variation" (CV), a measure that captures well human perception of the variation in a set of values (e.g., Weber, Shafir & Blais, 2004). The mean CV in the opinion-dependent condition was nearly one-fourth of that in the independent condition (0.15 vs 0.55), indicating significantly greater consensus among the opinions presented in the interdependent condition, $t(43)=34.1, p<.005$.

Results

Accuracy gains. Table 2 summarizes the mean absolute errors of the initial and final estimates and the accuracy gains in each condition. An overall 2x2 analysis of variance on the mean absolute errors showed a main effect of revision (phase 1 vs 2), $F(1,43)=14.13, p<.001$, no overall effect of sampling, $F(1,43)<1$, and, importantly, a significant interaction, $F(1,43)=19.54, p<.001$, suggesting that the change in accuracy was greater in the independent condition. Accuracy gains were obtained in both the independent condition, 27%, $t(43)=4.22, p<.001, d=0.64$, and the dependent condition, 7%, $t(43)=2.13, p<.05, d=0.32$. The gain was greater in the independent condition, $t(43)=4.42, p<.001, d=0.67$.

Frequency of revision. Table 2 also shows the percentage of times that participants gave final estimates that were *different* from their initial estimates. Participants changed their initial estimates more often in the independent than in the dependent condition, $t(43)=11.53, p<.001, d=1.7$.

Confidence. An overall 2x2 analysis of variance on the mean confidence showed main effects of revision (phase 1 vs 2), $F(1,43)=10.21, p<.001$, and sampling, $F(1,43)=15.64, p<.001$, and a significant interaction, $F(1,43)=16.55, p<.001$. Receiving advice increased participants' confidence in the dependent condition (59.4 vs 65.1), $t(43)=4.30, p<.001, d=0.65$, but not in the independent condition (58.5 vs 60.1), $p>.15$. Participants indicated greater confidence in their final estimates in the opinion-dependent than in the independent condition $t(43)=4.67, p<.001, d=0.7$.

Table 2
Results of Experiment

<u>Measure</u>	<u>Method of sampling advice</u>	
	Independent	Opinion-dependent
<u>Accuracy:</u>		
Initial error	100.4	93.9
Final error	73.1	87.2
% improvement	27	7
<u>Confidence:</u>		
Initial confidence	58.5	59.4
Final confidence	60.1	65.1
Increase in confidence	1.6	5.7
Rate of betting	.42	.58
<u>Revision Process:</u>		
Rate of changing initial estimates (%)	70	35

Betting. In accord with the confidence results, the participants bet more often in the dependent (58%) than in the independent condition (42%). The proportion of bets in the dependent condition (.58) was significantly greater than the chance rate of .5, $t(43)=3.76$, $p<.001$, $d=0.57$. This betting pattern supports our conclusion that participants trusted the consensus advice more than the disagreeing advice, despite being told about the sampling method. Had participants placed *all* their bets in the independent condition they would have earned a bonus at a rate of 24% compared with a rate of 17.5% in the dependent condition, $t(43)=2.46$, $p<.05$, $d=0.37$.

Conclusions. The results of the study reveal a confidence-accuracy dissociation, whereby participants were more confident in their less accurate judgments. We establish the dissociation using both confidence ratings and betting decisions. Two factors led to this dissociation effect. Participants offered higher confidence ratings and were more likely to bet on their estimates after receiving *opinion-dependent advice* since such advice involved greater consensus. Yet they gained more from *independent advice* than from redundant advice. This dissociation effect was obtained even though the participants were specifically informed about how the advisory opinions were selected on each trial. The participants revised their initial opinions more often in the independent condition, yet they apparently had limited trust in their revision results. Comparisons of the confidence ratings in phases 1 and 2 suggest that the conflicting advice (independent condition) did not affect the participants' confidence, but the agreeing advice (dependent condition) boosted it.

Discussion

Consensus is appealing and confidence-inducing, and for good reasons. A decision-maker (e.g., an editor or employer) acting on a set of opinions (referee reports or recommendations) is likely to feel more confident when the advisors appear to be in consensus than when they are in disagreement. Cialdini (1993, p. 103) notes that seeking and maintaining consistency in general is an adaptive rule of behavior, “a hallmark of logic and intellectual strength,” yet he warns of the risks of mindless endorsement of consistency. In the present research, consensus is a valid cue for confidence to the extent that the sources are known to be independent of each other. In contrast, spurious consensus (defined as a set of consistent opinions produced by interdependent sources) leads to dysfunctional judgment, specifically to dissociation between confidence and accuracy.

We used a belief-updating task in which the primary factor was the method of sampling the advisory opinions presented to participants for updating their beliefs. The method of sampling determined the interdependence and consensus levels and both of these factors affected the participants’ accuracy and confidence. Our participants reduced their errors by 27% upon receiving independent advisory opinions. This replicates a robust finding in the literature on advice-taking, whereby people improve their judgment accuracy appreciably merely by consulting randomly-drawn opinion(s) from a pool of estimates produced by others (Yaniv, 2004a; 2004b; Yaniv & Milyavsky, 2007). Participants gained less (7%) from using opinion-dependent advice. Nevertheless, they indicated higher confidence and showed a greater tendency to bet on the judgments that they had revised on the basis of opinion-dependent advice. As a result, we observed a dissociation between confidence and accuracy, whereby participants indicated greater confidence in and also a greater tendency to bet on judgments that were poorer overall.

In principle, participants’ confidence in advice-based judgments should be based on factors such as their own knowledge, their perceptions of the advisors’ expertise (Yaniv & Kleinberger, 2000), the consensus among the advisors, and the interdependence among them. As it turns out, participants pay close attention to consensus cues, that is, the level of agreement among the advisory opinions available to them, but not to the information on how the advisory opinions were sampled. Participants’ confidence ratings made before and after receiving advice confirm this conclusion. Obtaining opinion-dependent (and hence consensual) advice *increased* confidence, whereas receiving independently-drawn advice did not have such an effect, although the latter was more beneficial. Since independent opinions tended to be more conflicting, their value was not apparent to participants (cf. Budescu & Yu, 2006; Larrick & Soll, 2006).

Adaptive decision-making entails that factors that increase (decrease) the accuracy or validity of judgment should also increase (decrease) one's confidence. Why do people fail to appreciate the detrimental consequences of interdependence among sources, that is, why do they not lower their confidence accordingly? A cognitive analysis of the process of *learning from experience* offers a possible explanation that hinges on the perceptual asymmetry between consensus and interdependence.

Whereas consensus of opinion is typically quite salient, the amount of interdependence among sources of advice is not. To assess interdependence one would need to compute the correlations among advisors' judgments across cases or keep a record of the relationships among the advisors, their tendencies towards compliance, and the presence of sanctions or rewards that might influence their opinions. In a classic study investigating the evolution of beliefs among residents of MIT dormitories, Festinger, Schachter, and Back (1950) found that the group influence was associated with the underlying social network in the group, in which frequency of contact played an important role. In general, individuals are connected in a variety of networks formed by geographical proximity (neighbors, coworkers), social relationships (friendship, club membership), hierarchical relationships (position in an organization), and electronic contacts (internet). The intricacies of such networks presumably encumber decision-makers' ability to learn the role of interdependence and the differences between valid and spurious consensus. It is no wonder, then, that they fail to adjust for possible dependencies among the opinions they receive (cf. DeMarzo, Vayanos, & Zwiebel, 2003).

Relationship to Other Research

A related dissociation known from the literature on memory suggests that merely repeating an assertion leads people to judge it as more true or valid, regardless of the actual truth or falsity of the assertion (Hasher, Goldstein, & Toppino, 1977). The illusory truth of repeated statements is based on the fact that familiarity builds up automatically. Since people fail to monitor the source of the statements' familiarity, they are influenced by it in making their truth judgments (Begg, Anas, & Farinacci, 1992). Also related to the spurious consensus phenomenon is the confirmation bias, namely, people's tendency to attend selectively to confirmatory evidence, and their failure to appreciate the detrimental effects of biased assimilation of information (e.g., Lord et al., 1979).

The dissociation between confidence and accuracy is consistent with key findings in the literature on small groups. First, individuals who encounter disagreement among the other group members are less confident in their own opinion than those who encounter agreement. Second, there is evidence that disagreement among group members is associated with greater accuracy

gains and, moreover, heterogeneous groups exhibit better performance than homogeneous ones (Mannix & Neale, 2005; Snizek, 1992; Sommers, 2006). The groupthink phenomenon suggests that social pressures in homogeneous groups lead to spurious consensus. Related to that is also the link made more recently by Sunstein (2003), in his treatise “Why society needs dissent,” between consistency pressures (e.g., in the legal system) and “informational cascades” which could sometimes, though not always, lead to mistaken “herd behavior” (Bikhchandani, Hirshleifer, & Welch, 1992).

Limitations and Reservations

Our findings need to be qualified in several ways. First, although interdependence has been at the focus of this research, not all forms of statistical interdependence are of interest to psychologists; in fact, most are not. For example, two sources A and B would be interdependent if A were to double (or to add a constant to) the numerical estimates given by B. Such functional forms of interdependence are of little relevance here, since they do not correspond to interdependence patterns that could conceivably be created via commonly occurring social interactions. We have focused solely on one form of interdependence, namely, the form resulting from common mechanisms of social influence whereby person B comes to express an opinion similar to that of person A.

Second, we have considered the use of opinions on matters of fact, that is, opinions whose accuracy can be verified objectively. Can these findings be readily extended to advice on matters of taste? We think that different rules guide the search for good advice on subjective issues. When people solicit others’ advice on a movie or a dish in a restaurant, their main objective is often to predict how they themselves will react to the same movie or dish when they experience it. With this goal in mind they may seek the opinions of similar individuals who have made correlated rather than independent sequences of choices in the past. The development of these ideas is challenging, but beyond the scope of the present paper.

Third, the effect of spurious consensus on confidence may also be limited. For instance, it may not hold in adversarial or strategic settings where decision-makers suspect that the advisors may be motivated to mislead them. For instance, shoppers are unlikely to have greater trust and confidence in salespeople who provide identical opinions on the virtues of a product on sale.

Fourth, the present experimental results demonstrate that, under ordinary circumstances, people may easily be led to express greater confidence in their less accurate decisions. We do not claim that confidence and accuracy are generally dissociated. In fact, confidence judgments are often found to be monotonically related to accuracy level, even if they are miscalibrated.

Finally, we do not claim that our participants were completely oblivious to the importance of independence of sources of information; they may have been partially sensitive to such information, albeit insufficiently so. We only claim (a) that the dissociation between confidence and accuracy occurs when interdependence and consensus cues have *opposite* or *conflicting* influences on the decision-maker's confidence and, (b) that such conflicting influences arise due to spurious consensus. Decision-makers faced with selections of consensual opinions (or pieces of evidence) need to consider the selection methods used to create these sets. A failure to assess the detrimental effects of spurious consensus could lead to the dissociation between confidence and accuracy witnessed in this research. Future research might investigate under what conditions participants utilize information about source interdependence and thus avoid such pitfalls.

References

- Begg, I., Anas, A., & Farinacci, S. (1992). Dissociation of processes in belief: Source recollection, statement familiarity, and the illusion of truth. *Journal of Experimental Psychology: General*, 121, 446-458.
- Bikhchandani, S., Hirshleifer, D., & Welch, I. (1992). A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of Political Economy*, 100, 992-1026.
- Budescu, D.V., & Yu, H.-T. (2006). To Bayes or not to Bayes: A comparison of two classes of models of information aggregation. *Decision Analysis*, 3, 145-162.
- Budescu, D.V., & Yu, H.-T. (2007). Aggregation of opinions based on correlated cues and advisors. *Journal of Behavioral Decision Making*, 20, 153-177.
- Cialdini, R.B. (1993). *Influence: The psychology of persuasion*. New York: William Morrow.
- DeMarzo, P.M., Vayanos, D., & Zwiebel, J. (2003). Persuasion bias, social influence and uni-dimensional opinions. *Quarterly Journal of Economics*, 118, 909-968.
- Festinger, L., Schachter, S., & Back, K. (1950). *Social pressures in informal groups: A study in human factors in housing*. Stanford: Stanford University Press.
- Harries, C., Yaniv, I., & Harvey, N. (2004). Combining advice: The weight of a dissenting opinion in the consensus. *Journal of Behavioral Decision Making*, 17, 333-348.
- Harvey, N., & Fischer, I. (1997). Taking advice: Accepting help, improving judgment and sharing responsibility. *Organizational Behavior and Human Decision Processes*, 70, 117-133.
- Hasher, L., Goldstein, D., & Toppino, T. (1977). Frequency and the conference of referential validity. *Journal of Verbal Learning and Verbal Behavior*, 16, 107-112.
- Hogarth, R.M. (1978). A note on aggregating opinions. *Organizational Behavior and Human Performance*, 21, 40-46.
- 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.
- Kahneman, D., & Tversky, A. (1973). On the psychology of prediction. *Psychological Review*, 80, 237-251.
- Kroll, Y., Levy, H., & Rapoport, A. (1988). Experimental tests of the separation theorem and the capital asset pricing model. *American Economic Review*, 78, 500-519.
- Larrick, R.P., & Soll, J.B. (2006). Intuitions about combining opinions: Misappreciation of the averaging principle. *Management Science*, 52, 111-127.
- Lord, C.G., Ross, L., & Lepper, M.R. (1979). Biased assimilation and attitude polarization: The effects of prior theories on subsequently considered evidence. *Journal of Personality and Social Psychology*, 37, 2098-2109.
- Maines, L. (1996). An experimental examination of subjective forecast combination. *International Journal of Forecasting*, 12, 223-233.
- Mannix, E., & Neale, M.A. (2005). What differences make a difference? The promise and reality of diverse teams in organizations. *Psychological Science in the Public Interest*, 6, 31-55.
- Savadori, L., Van Swol, L.M., & Snizek, J.A. (2003). Information sampling and confidence within groups and judge advisor systems. *Communication Research*, 28, 737-771.

- Snizek, J.A. (1992). Groups under uncertainty: An examination of confidence in group decision making. *Organizational Behavior and Human Decision Processes*, 52, 124-155.
- Soll, J.B. (1999). Intuitive theories of information: Beliefs about the value of redundancy. *Cognitive Psychology*, 38, 317-346.
- Sommers, S.R. (2006). On racial diversity and group decision making: Identifying multiple effects of racial composition on jury deliberations. *Journal of Personality and Social Psychology*, 90, 597-612.
- Sunstein, C.R. (2003). *Why societies need dissent*. Cambridge, MA: Harvard University Press.
- Weber, E.U., Bockenholt, U., Hilton, D.J., & Wallace, B. (2000). Confidence judgments as expressions of experienced decision conflict. *Risk, Decision and Policy*, 5, 69-100.
- Weber, E.U., Shafir, S., & Blais, A.-R. (2004). Predicting risk sensitivity in humans and lower animals: Risk as variance or coefficient of variation. *Psychological Review*, 111, 430-445.
- Yaniv, I. (1997). Weighting and trimming: Heuristics for aggregating judgments under uncertainty. *Organizational Behavior and Human Decision Processes*, 69, 237-249.
- Yaniv, I. (2004a). The benefit of additional opinions. *Current Directions in Psychological Science*, 13, 75-78.
- Yaniv, I. (2004b). Receiving other people's advice: Influence and benefit. *Organizational Behavior and Human Decision Processes*, 93, 1-13.
- Yaniv, I., & Kleinberger, E. (2000). Advice taking in decision making: Egocentric discounting and reputation formation. *Organizational Behavior and Human Decision Processes*, 83, 260-281.
- Yaniv, I., & Milyavsky, M. (2007). Using advice from multiple sources to revise and improve judgment. *Organizational Behavior and Human Decision Processes*, 103, 104-120.