Journal of Behavioral Decision Making, *J. Behav. Dec. Making* (2011)
Published online in Wiley Online Library (wileyonlinelibrary.com) **DOI:** 10.1002/bdm.740

Exploiting the Wisdom of Others to Make Better Decisions: Suspending Judgment Reduces Egocentrism and Increases Accuracy

ILAN YANIV* and SHOHAM CHOSHEN-HILLEL
Hebrew University of Jerusalem, Jerusalem, Israel

ABSTRACT

Although decision makers often consult other people's opinions to improve their decisions, they fail to do so optimally. One main obstacle to incorporating others' opinions efficiently is one's own opinion. We theorize that decision makers could improve their performance by suspending their own judgment. In three studies, participants used others' opinions to estimate uncertain quantities (the caloric value of foods). In the full-view condition, participants could form independent estimates prior to receiving others' opinions, whereas participants in the blindfold condition could not form prior opinions. We obtained an intriguing blindfold effect. In all studies, the blindfolded participants provided more accurate estimates than did the full-view participants. Several policy-capturing measures indicated that the advantage of the blindfolded participants was due to their unbiased weighting of others' opinions. The full-view participants, in contrast, adhered to their prior opinion and thus failed to exploit the information contained in others' opinions. Moreover, in all three studies, the blindfolded participants were not cognizant of their advantage and expressed less confidence in their estimates than did the full-view participants. The results are discussed in relation to theories of opinion revision and group decision making. Copyright © 2011 John Wiley & Sons, Ltd.

KEY WORDS judgment; decision making; advice taking; combining opinions; belief revision

It is a common practice to solicit other people's opinions prior to making a decision. A student seeks other students' ratings of an elective course, and a manager considers several judgmental forecasts of foreign exchange rates before embarking on a new venture. Such settings involve the decision maker in the task of combining other people's opinions, mostly to improve one's final decision (Gino, Shang, & Croson, 2009; Harvey & Fischer, 1997; Yaniv, Choshen-Hillel, & Milyavsky, 2011).

In some settings, decision makers form their own preliminary views prior to soliciting others' opinions; thus, their task, upon receiving the advice, is to revise their opinions. In other settings, decision makers approach the task *tabula rasa*, that is, with hardly any prior knowledge that would enable them to form clear opinions; here, their task is merely to combine others' advice. For example, a student (or a manager) seeking others' opinions might or might not have the information to form a confident, independent preliminary forecast.

How might the presence or the absence of prior opinions affect decision makers' ability to use advice profitably? Do prior opinions aid one's performance (e.g. by adding information) or hamper it? How might suspending judgment affect accuracy? We suggest that judges engage different modes of processing when integrating others' opinions, depending on whether or not they hold a prior opinion of their own. These modes, which could be traced using process measures, determine the judges' success in judgmental estimation tasks.

Specifically, we suggest that judges who do *not* hold personal opinions form an aggregate opinion by attending to all opinions and assessing an intuitive measure of the central tendency in the set (Budescu, Rantilla, Yu, & Karelitz, 2003; Budescu & Yu, 2007). Normative studies have shown that statistical equal weighting of judgments yields aggregate forecasts that are more accurate than the individual opinions on which they are based (Larrick & Soll, 2006). In the case of quantitative judgments, we can outline in simple terms why improvement should occur when estimates are combined. A subjective estimate of an objective event can be viewed as the sum of three components: the "truth," random error (random fluctuations in a judge's performance), and constant bias (a consistent tendency to overestimate or underestimate the event). Statistical principles guarantee that judgments formed by averaging several sources have lower random error than the individual sources on which the averages are based. Therefore, if the bias is small or zero, the average judgment should converge about the truth (Einhorn, Hogarth, & Klempner, 1977). Indeed, accuracy gains from combining opinions have been observed in a variety of domains, ranging from judgments of physical quantities to forecasts of business outcomes (Harvey & Fischer, 1997; Rowe & Wright, 2001; Soll & Larrick, 2009; Surowiecki, 2004). Thus, to the extent that participants approximate normative equal weighting, they should improve their accuracy.

In contrast, we suggest that judges who *hold* prior opinions engage in a different kind of revision process. Their prior beliefs bias their weighting of others' advice so that in combining their own opinions with those of others, they assign proportionally less weight to others' opinions than to their own. Although judges holding prior opinions improve the accuracy of their estimations by consulting additional opinions, their gains are suboptimal because they fail to fully exploit the "wisdom of others"—the information contained in others' opinions (Harvey & Harries, 2004; Mannes, 2009; Yaniv & Milyavsky, 2007). The accuracy costs of this self-other bias are considerable. They become even greater when *several* other opinions are available, because participants assign the other opinions about the same weight altogether,

*Correspondence to: Ilan Yaniv, Department of Psychology & Center for the Study of Rationality, Hebrew University of Jerusalem, Jerusalem 91905, Israel.
E-mail: ilan.yaniv@huji.ac.il

almost irrespective of their number (e.g. two, four, or eight, Yaniv & Milyavsky, 2007).

One cognitive account of such egocentric discounting of others' opinions is that individuals are privy to their own thoughts but not to those of others. If we suppose that people have more access to evidence supporting their own views than to evidence supporting others' views and that their weighting of opinions is a function of the evidence available, then it follows that they should assign more weight to their own than to others' opinions (Yaniv, 2004a; see review in Bonaccio & Dalal, 2006).

That knowledge updating may be held back by prior conceptions has been amply shown in studies in a variety of domains (Tetlock, 2005). Generally, individuals tend to be overly conservative in changing their prior hypotheses (Klayman & Ha, 1987). Research on intuitive hypothesis testing has shown that judges tend to seek information that agrees (rather than disagrees) with their prior hypotheses. Students of attitude change have also documented people's tendency to persist in their attitudes, because of their biased weighting of evidence (Cohen, Aronson, & Steele, 2000; Lord, Ross, & Lepper, 1979). In sum, the findings in the area of advice taking (e.g. Harvey & Harries, 2004) are consistent with the general conclusion that one's incorporation of new evidence is biased in the direction of one's prior beliefs.

Our studies manipulated decision makers' ability to generate prior opinions. The participants who were assigned to the *full-view* condition could generate initial opinions (estimates of the caloric values of various foods), whereas the ones assigned to the *blindfold* condition could not. We evaluated participants' policies for weighting other opinions, their judgmental accuracy, and their confidence. The full-view participants were expected to display an egocentric bias in updating their prior opinions on the basis of advice. The blindfolded participants (who were forced to suspend their personal opinions), in contrast, were expected to weight other people's opinions more equally and thus more profitably. We thus expected the participants who did *not* form prior opinions to gain more from the advice presented to them. Although we hypothesized that the blindfolded participants would perform better than the full-view participants, we also expected that they would paradoxically feel less confident (compared with the full-view participants), as they would regard their inability to form prior opinions as a disadvantage.

STUDY 1

Method

The experimental procedure was conducted individually on personal computers and included 25 questions on the caloric value of various foods (e.g. "the number of calories in a bowl of cooked rice"). Participants ($n = 67$) were undergraduate students. They were told that they would get a flat fee of 7 Israeli Shekels (IS) and 1 IS (about $0.25) for each estimate that fell within the range extending 15% on either side of the correct answer.

There were two between-participants conditions (full view versus blindfold). The *full-view* condition began with a practice part that included five trials intended to familiarize the participants with the types of questions used. On each trial, the participants were presented with a question and requested to enter their opinion. No advice was given during practice, and the data from these trials were not analyzed. Then, for the main part, 20 questions were presented. For each question, the participants' initial opinion (number of calories in the target food) was elicited. Then, after they had entered their best estimates, five additional opinions were displayed, one below the other (Table 1). Thus, the initial estimate was listed first, followed by five advisory estimates. These estimates were drawn from pools that each included 100 estimates collected in an earlier survey.

The advisory estimates were sampled at random, with labels such as #6, #10, and #29, indicating that the estimates came from different individuals on each trial. The participants were specifically told that the opinions had been drawn *at random* by the computer from large pools of estimates made by other participants. After viewing the advisory estimates, the participants were asked to make their final, possibly revised, caloric estimates. In addition, they were asked to rate their confidence that their estimates fell within the range extending 15% on either side of the correct answer; the scale was anchored at 0% (not confident at all) and 100% (completely confident), with a tick point at every 10%. This procedure was repeated for each of the 20 calorie questions.

The *blindfold* condition also began with five practice questions for which the participants were requested to enter their opinions. Then, for the main part, the participants were told that they would be presented with 20 calorie questions, except that the name of the target food in each would be replaced by a randomly selected letter code, for example, "the number of calories in one serving of K." Further, they were told that they would be provided with six advisory estimates made by other participants *who knew the names of the target foods* (Table 1). On each trial, the participants were asked to estimate the caloric value of the (concealed) target food and indicate their confidence in their estimate.

Table 1. Sample materials in Study 1

Full-view condition	
What is the calorie value of an orange?	
Your best estimate was	100
The best estimate of advisor #26	90
The best estimate of advisor #4	84
The best estimate of advisor #19	320
The best estimate of advisor #97	140
The best estimate of advisor #12	50
Your final best estimate	_____
Blindfold condition	
What is the calorie value of "D"?	
The best estimate of advisor #11	120
The best estimate of advisor #26	90
The best estimate of advisor #4	84
The best estimate of advisor #19	320
The best estimate of advisor #97	140
The best estimate of advisor #12	50
Your final best estimate	_____

DOI: 10.1002/bdm

In the full-view condition, the first-listed estimate on each trial was the participant's *initial own* estimate, whereas in the blindfold condition, the first-listed estimate was an advisory estimate sampled *at random*.

Results

The mean absolute errors of the participants' estimates (Table 2) served as a measure of accuracy. In the full-view condition, the participants' final estimates were more accurate than their initial estimates (76.2 and 91.2), $t(33) = 7.49$, $p < 0.01$, $d = 1.29$, thereby replicating the standard finding that receiving advice is beneficial. Importantly, the participants' final estimates were even more accurate in the blindfold than in the full-view condition (66.0 vs 76.2), $t(65) = 2.55$, $p < 0.05$, $d = 0.64$.

Do the full-view participants perform less well because they discount advice? The evidence of egocentric discounting of advice comes from two measures. First, the full-view participants made final judgments that were *identical* to their initial ones in 52.5% of the cases, whereas the blindfolded participants adopted their first-listed estimates in only 12.3% of the cases, $t(65) = 9.10$, $p < 0.001$, $d = 2.35$. Second, the final estimates of the full-view judges were closer to their initial estimates (mean distance = 26.6) than were those of the blindfolded judges to their first-listed advisory estimates (91.3), $t(65) = 14.16$, $p < 0.001$, $d = 3.49$. These results provide evidence of egocentric discounting in the full-view condition.

Moreover, the blindfolded participants weighted the advisory estimates more evenly. The distance of the participant's final estimate from each of the six opinions indicates how much weight each opinion received. For the full-view participants, these distances averaged 26.6, 102.0, 99.2, 105.9, 99.8, and 95.5 (1st through 6th estimates, respectively). For the blindfolded participants, the distances averaged 91.3, 84.4, 84.7, 88.8, 86.8, and 83.8.

Subsequent analyses, shown in Table 2, further corroborate the idea that egocentric bias is detrimental to accuracy. Averaging the six opinions produced more accurate estimates than did either the full-view judges (55.2 vs 76.2), $t(33) = 6.90$, $p < 0.001$, $d = 1.18$ or the blindfolded judges (54.9 vs 66.0), $t(32) = 6.10$, $p < 0.001$, $d = 1.06$. The superior accuracy of equal weighting over the judges' idiosyncratic weighting policies suggests that judges (in both conditions) underutilized some of the information contained in the opinions available to them. The blindfolded judges departed from equal weighting to a lesser extent than did the full-view ones. Specifically, the distance of a final estimate from the *simple average* of the six estimates (presented on the same trial) was shorter in the blindfold condition than in the full-view condition (37.9 vs 53.8), $t(65) = 5.50$, $p < 0.001$, $d = 1.36$. The greater proximity of the blindfolded judges to equal weighting presumably accounts for their greater accuracy.

Finally, although their final estimates were more accurate, the blindfolded participants were less confident in them than the full-view participants (47 vs 56%), $t(65) = 2.23$, $p < 0.05$, $d = 0.55$. Presumably, they felt that not knowing the names of the target foods was a disadvantage in the estimation task.

Discussion

The full-view participants, who were asked to state their initial opinions prior to observing others' advice, revised their opinions in a biased manner. Their final judgments were *less* accurate, presumably because they discounted the advisory opinions and failed to exploit the information contained in them. The blindfolded participants, who could not generate initial opinions, weighted the other opinions more equally and thereby achieved greater accuracy.

We conclude that the full-view participants wasted some of the information contained in the opinions. The extent of this waste could be assessed by a simulation analysis comparing the accuracy of the full-view participants with that of simple averages of randomly drawn samples of opinions. To this end, samples of opinions of size k were drawn, where k varied from one to six. The average distance of *single* opinions from the truth was 91.2. The average distance of the averages of *pairs* of opinions was 74.1, about the level of accuracy achieved by the full-view participants (76.2) who had a total of *six* opinions available to them. The accuracies (average distances) of opinion samples of sizes three, four, five, and six were 66.8, 62.5, 58.8, and 55.2, respectively, better than that achieved by the full-view participants.

Where exactly does the disadvantage of the full-view protocol lie? This protocol differed from the blindfold

Table 2. Results of Studies 1–2

	Study 1		Study 2	
	Full view	Blindfold	Full view	Blindfold
	($n = 34$)	($n = 33$)	($n = 45$)	($n = 44$)
Accuracy (mean absolute errors)				
Initial estimate	91.2	—	—	—
Final estimate	76.2	66.0	76.0	66.1
Average (equal weight) of six estimates	55.2	54.9	53.6	54.8
Measures of egocentrism				
% keeping the first-listed estimate[a]	52.5	12.3	—	—
Distance between final and first-listed estimates[a]	26.6	91.3	—	—
Distance between final estimate and average of six estimates	53.8	37.9	69.4	43.4
Confidence				
Rating (0–100% scale)	56	47	56	48

[a]The first-listed estimate is the participant's initial estimate in the full-view condition and a randomly drawn estimate in the blindfold condition.

DOI: 10.1002/bdm

protocol in two aspects. First, critical information (target food name) was available to the participants, and, second, their initial estimates were elicited. Are both of these aspects necessary for the effects to occur?

According to one plausible hypothesis, an important driver of the egocentric effect is the explicit elicitation of initial estimates. Once this factor is removed, performance in the full-view condition should no longer be inferior to that in the blindfold condition. This hypothesis presupposes that judges do not create estimates unless they are asked to. According to the second hypothesis, merely knowing the name of the food suffices to create an egocentric effect. This is because judges form internal estimates, even when they are not asked to do so, and they weight the advisory opinions preferentially, giving more weight to the ones near their own internal estimates. Study 2 was designed to test these two possibilities.

STUDY 2

The participants in the blindfold and the full-view conditions were shown six advisory estimates on each trial, based on which they were supposed to estimate the caloric value of some food. The full-view participants were shown the name of the food, whereas the blindfolded participants were shown an arbitrary letter code substituting for the food name. Unlike in Study 1, *no* initial estimates were elicited in the full-view condition. The two conditions thus differed only in that the food name was shown in one but not in the other.

Method

As in Study 1, the procedure was conducted individually on personal computers and included 25 questions on the caloric value of various foods. Participants ($n = 89$) were told that they would get a bonus of 1 IS (about \$0.25) for each estimate close to the truth, in addition to their flat fee of 7 IS.

There were two between-participants conditions (full view versus blindfold). As in Study 1, the *full-view* condition began with a practice part in which the participants were simply requested to estimate the caloric value of five foods. Then, for the main part, 20 questions were presented, each one along with six advisory opinions (drawn from relevant pools of estimates). Thus, as opposed to Study 1, the participants were only asked to make *final* estimates (initial estimates were not elicited). After entering the final estimates, the participants were asked to rate their confidence that their estimates fell within the range extending 15% on either side of the correct answer, on a 0–100% scale. This procedure was repeated for each of the 20 calorie questions.

The *blindfold* condition was identical to that in Study 1. Thus, it differed from the present full-view condition only in that the target food in each question was concealed. Following the practice part (with five questions), the participants were told that they would be presented with 20 calorie questions, except that the names of the foods would be replaced by randomly selected letter codes. In addition, they were told that they would be provided with six advisory estimates made by other participants *who knew the names of the target foods*. On each trial, the participants were asked to estimate the caloric value of the (concealed) target food and indicate their confidence in their estimate.

Results

The pattern of results of Study 2 closely resembles that of Study 1 (Table 2). The blindfolded judges were more accurate than the full-view judges (mean absolute errors: 66.1 vs 76.0), $t(87) = 2.18$, $p < 0.05$, $d = 0.47$. Thus, the full-view judges were less accurate simply because they knew the names of the target foods, and thus, they used more idiosyncratic (rather than equal) weighting policies. The conjunction of the following two facts corroborates this interpretation. First, the equal-weighting policy (averaging of the six opinions) produced more accurate estimates than did either the full-view judges (53.6 vs 76.0), $t(44) = 5.98$, $p < 0.001$, $d = 0.89$ or the blindfolded judges (54.8 vs 66.1), $t(43) = 5.36$, $p < 0.001$, $d = 0.81$. Second, the blindfolded participants used judgmental weighting policies that were more similar to equal weighting than those that the full-view participants used. Specifically, the distance of a final estimate from the *simple average* of the six advisory estimates (presented on that trial) was shorter in the former than in the latter (43.4 vs 69.4), $t(87) = 6.06$, $p < 0.001$, $d = 1.30$.

Finally, confidence and accuracy were again dissociated. Although the blindfolded participants were more accurate, they felt less confident in their estimates than the full-view participants (48 vs 56%), $t(87) = 2.35$, $p < 0.05$, $d = 0.5$.

STUDY 3

In Studies 1–2, the full-view participants could form their personal opinions at an early stage of the advice-aggregation process. We have seen that such opinions bias the participants' processing of additional opinions and hinder their performance. Suppose, however, that the participants are allowed to form their personal opinions only *after* forming their advice-based, blindfold estimates. Might the personal opinions thus formed be less egocentric and more accurate? One possibility is that one's personal opinions always take priority, even when formed at a later stage. Another possibility is that personal opinions generated after forming a blindfold (advice based) estimate are less biased. By manipulating the timing of the formation of personal opinions, we tested the impact of suspending judgment on reducing egocentric discounting. Aside from their theoretical importance, the findings should have practical significance in suggesting ways of improving common practices for combining opinions.

Whereas the full-view condition remained the same as in Study 2, the blindfold condition was modified as follows. On each trial, the participants were asked to produce a blindfold (advice based) estimate. But then, the actual name of the target food was disclosed to them, and they could bring to bear their personal opinions and knowledge in producing their second (possibly revised) estimates.

DOI: 10.1002/bdm

Method

Ninety-two participants were assigned at random to one of two conditions. The full-view condition was similar to the one in Study 2, except that five (rather than six) advisory estimates were presented on each trial. The blindfold condition was also the same as in Study 2, except that it included two estimates of the calories in two phases. On each trial, the participants first entered their blindfold estimate of the number of calories based on five opinions (*blindfold phase*). Then, the name of the target food was disclosed, and the participants were prompted to produce a final estimate of the number of calories in the food (*disclosing phase*). Here, the participants could either keep or change their blindfold estimate, using their knowledge about the target food.

Results

As shown in Table 3, the blindfold estimates were more accurate than the full-view one (mean absolute errors: 63.9 vs 78.2), $t(90) = 3.52$, $p < 0.01$, $d = 0.74$. The estimates after the target food names were disclosed were also more accurate than those in the full-view condition (66.2 vs 78.2), $t(90) = 2.51$, $p < 0.05$, $d = 0.53$. Thus, finding out the name of the food at a later phase did not degrade the participants' performance (66.2 vs 63.9), $t < 1$. Notably, the participants' accuracy levels remained about the same, although they did *not* keep their blindfold estimates, but *changed* them in about 80% of the cases after finding out the name of the food (the average distance between the estimates given in the two phases was 55.8).

The blindfold estimates were closer to the average advice than were the full-view estimates (36.2 vs 69.4), $t(90) = 8.29$, $p < 0.001$, $d = 1.75$ and so were the estimates made after the food names were disclosed (58.3 vs 69.4), $t(90) = 2.46$, $p < 0.05$, $d = 0.52$. Thus, in forming estimates in the blindfold phase *and* in the disclosing phase, the participants tended more toward equal weighting of all opinions and hence benefited more from the advice than did the full-view participants.

Finally, a confidence–accuracy dissociation was found as in previous studies. The blindfolded participants (in the blindfold phase) expressed lower confidence than did the full-view participants (44 vs 58%), $t(90) = 3.78$, $p < 0.001$, $d = 0.80$, even though they performed better. The blindfolded participants were presumably not cognizant of the advantage conferred by the blindfolding manipulation; their confidence was boosted significantly once the target names were revealed (44 vs 60%), $t(45) = 7.74$, $p < 0.001$, $d = 1.14$, although finding out the food names did not contribute to their accuracy. According to our analyses from Study 1, the marginal contribution (to accuracy) of a sixth opinion is minute. Our participants felt, though, that their personal opinions (sixth in order) contributed significantly to their accuracy.

GENERAL DISCUSSION

Under what conditions do people receive the full benefit of the "wisdom of others"? Our findings suggest that decision makers' prior opinions *hamper* their ability to use advice to its full potential. Thus, contrary to what one might expect, decision makers achieve more by approaching a task *tabula rasa*—that is, without prior opinion or knowledge. Three patterns of findings are particularly informative here—first, the effects of blindfolding on accuracy; second, the *dissociation* between accuracy and confidence; and third, the judges' policies for using advice.

Blindfolding effects on accuracy

In three studies, our participants were asked to make decisions based on other people's advisory opinions. We compared two conditions. In one, the participants could form prior estimates (full-view condition), whereas in the other, they could not form their own personal opinions but had to rely only on advice (blindfold condition). In all studies, the blindfolded judges used the advice better and gained more in accuracy than did the full-view judges. The blindfolding effect on accuracy was obtained in Studies 1 and 2, despite an important difference between them—namely, the participants were asked to record their initial opinions in Study 1 but not in Study 2. In Study 3, the participants in the blindfold condition first gave their blindfold judgments and then, after seeing the target food name, their revised judgments. Importantly, their revised estimates were still more accurate than the estimates given by the full-view participants. This result supports the hypothesis that suspending judges' ability to form personal opinions until after they had produced blindfold (i.e. advice based) estimates enhances their use of the advice and their final accuracy.

Table 3. Results of Study 3

	Full view	Blindfold	
	($n = 46$)	($n = 46$)	
		Blindfold phase	Disclosing phase
Accuracy (mean absolute errors)			
Estimate	78.2	63.9	66.2
Average (equal weight) of five estimates	58.3	56.6	56.6[a]
Measures of egocentrism			
Distance between estimate and average of five estimates	69.4	36.2	58.3
Confidence			
Rating (0–100% scale)	58	44	60

[a]The average shown in the blindfold phase also applies here.

DOI: 10.1002/bdm

Confidence–accuracy dissociations

Our analysis of judges' confidence in their estimates in relation to their accuracy sheds further light on their processing of advisory opinions. Accuracy and confidence were dissociated in all studies. Although the blindfolded participants were, on average, more accurate than the full-view participants, they felt less confident in their estimates. Such dissociations are of interest because they contrast with the common finding of a positive correlation between confidence and accuracy (e.g. Ronis & Yates, 1987). The confidence–accuracy correlation broke down in our case, presumably because the blindfolded participants failed to realize that "not knowing the name of the food" was to their advantage. Likewise, our full-view participants failed to realize that forming prior opinions had a negative effect on their ability to integrate others' opinions fully. The findings of Study 3 further strengthen this interpretation; the blindfolded participants' confidence was boosted unduly when the food names were disclosed, although their personal opinion had negligible marginal value when added to the five advisory opinions already available.

Judges' policies in using advice

In principle, judges' success depends on their policies in using advisory opinions. We have theorized that judges who hold prior opinions benefit less from receiving others' advice, because they persist in their prior opinions and discount those of others (Cohen et al., 2000; Lord et al., 1979). In contrast, judges who do not hold prior opinions benefit more from receiving others' opinions, because they tend to give more equal consideration to all opinions and thus utilize the information more effectively.

Our findings indeed show that the full-view and blindfold conditions triggered different modes of processing others' opinions. Our full-view participants showed an egocentric bias in the revision process. They adhered to their prior opinions in as many as 52.5% of the cases, whereas in the remaining cases, they made final estimates that were close to their initial estimates (Study 1). The blindfolded participants tended to weight the opinions more equally—their final estimates were closer to the average of the advisory opinions than were those made by their full-view counterparts. Our conclusion here—that individuals rely on different modes of processing advice—is remarkably similar to that reached by Soll and Mannes (2011).

Why might judges persist in their opinions?

Why do judges overweight their own opinions and fail to take full advantage of the other opinions? Our explanation is that judges have differential-access to evidence supporting their own and others' opinions. In the process of forming their final opinions, judges summarize their relevant internal knowledge base. Naturally, they are privy to the reasons supporting their own estimates as well as the strength of those reasons but are not privy to the internal network of reasons underlying the advisors' opinions. Therefore, they discount opinions for which they have less justification (Yaniv, 2004a). The results of our studies are consistent with this account. In Study 1, the full-view participants formed an explicit opinion and assigned more weight to it than to alternative advisory opinions. In Study 2, the full-view participants were *not* asked to enter their initial opinion. The data imply that the participants still found some opinions more appealing than others and, as a result, did not weight all opinions equally. Study 3 tested further how the timing of forming one's personal opinion might affect processing. The participants in the blindfold condition who were told the name of the target food after they had formed their blindfolded opinions also tended toward equal weighting of the various opinions. Thus, suspending judges' ability to form personal opinions until after they had produced blindfold (i.e. advice based) estimates leads them to more egalitarian use of the advice.

Several alternative accounts of egocentric discounting could also be suggested (e.g. Bonaccio & Dalal, 2006), but they all seem less parsimonious, as they do not readily account for the whole range of past and present empirical results. Consider first the classic anchoring-and-adjustment heuristic (Tversky & Kahneman, 1974). The participants' initial estimates might have served as an anchor to their final estimates. In Studies 2 and 3, however, initial self opinions were not elicited, so no explicit anchors were available. Moreover, the participants were presented with multiple pieces of advice along with each question. Prominent theories of anchoring (e.g. Strack & Mussweiler, 1997) discuss the integration of a *single* initial value into one's final estimation and say nothing about the integration of *multiple* values. Thus, many new assumptions would have to be introduced to explain how the participants integrated the multiple pieces of advice presented to them in Studies 2–3.

According to another alternative account, judges' commitment to their prior opinions leads to the egocentric discounting of advice. Indeed, the motivation to maintain consistency could be powerful if a decision maker's opinion is to be made public or if a change of opinion is perceived as ego threatening, impinging on one's self-perception and deep-rooted political or religious beliefs. However, the conditions that typically create and enhance commitment to one's own position were not present in our procedure, which involved neutral questions, was conducted in isolation in front of a computer terminal, and provided rewards for accuracy. Moreover, an explanation in terms of commitment and consistency does not readily account for judges' sensitivity to the changing quality of advice (Yaniv & Kleinberger, 2000), their sensitivity to their own expertise (Yaniv, 2004b), and the opinion change at the disclosing phase in Study 3.

Finally, a third alternative account for judges' egocentric discounting is that they believed their opinions were superior to those of others. Harvey and Harries (2004) indeed report some data consistent with this explanation, as well as data that seem consistent with our account of advice taking. However, the "superiority account" by itself does not predict that judges' weighting policies should be sensitive to the

DOI: 10.1002/bdm

quality of the advice (Yaniv & Kleinberger, 2000) as well as to their own level of expertise (Yaniv, 2004b). In contrast, the differential-access explanation, namely, the idea that judges consider the evidence they can access for the various opinions and weight them accordingly — although in a biased manner — could explain judges' overweighting of their own opinion as well as the findings showing sensitivity to the knowledge factor.

Relationship to other research

Our findings on egocentric discounting echo Edwards's (1968) seminal work on *conservatism* in Bayesian updating, which used the classic bookbag and poker-chip task. In this task, participants sample probabilistic information and update their prior hypotheses online. Edwards's participants adjusted their initial hypotheses far too slowly, as if they underweighted the informational value of the samples obtained in the task. Edwards (1968) summarized his results by saying that "a convenient approximation to the data would say that it takes anywhere from two to five observations to do one observation's worth of work in inducing a subject to change his opinions" (p. 17). Our own documentation of informational *waste* in the discussion of Study 1 is consistent with Edwards' conclusion. It appears that judges adhere not only to their internally generated hypothesis or initial opinions, as in our studies, but also to externally generated ones, as in the classic updating task (see also Ronis & Yates, 1987).

Our findings also tie in with Tetlock's (2005, chaps. 3–4) study of the quality of experts' forecasts about economic and political affairs. Tetlock's main finding is that experts' open-mindedness is associated with their forecast accuracy and coherence. Experts who adhere to a well-defined worldview perform less well than those who are flexible in their thinking and willing to consider disconfirming evidence. Our experimental findings are consistent with this view in that the removal of the opportunity to form preconceptions benefits performance. In a different vein, Herzog and Hertwig's (2009) study of the "wisdom of many within one mind" also shows the utility of disconfirmation procedures designed to counteract one's fixation on one's prior opinions. Such procedures prove effective in eliciting multiple estimates from the same person (see also Vul & Pashler, 2008).

Finally, our findings on the disjoint effects of blindfolding on the participants' confidence and accuracy tie in with earlier reports of confidence–accuracy dissociation. The trigger for this dissociation in the present study was the blindfolding versus full-view manipulation. In a related study by Yaniv, Choshen-Hillel, and Milyavsky (2009), *spurious consensus* triggered the same type of dissociation. Thus, interdependent, consensual sources induced unduly high confidence, whereas independent sources led to greater accuracy and lower confidence. Hall, Ariss, and Todorov (2007) showed that illusion of knowledge could have divergent effects on accuracy and confidence. Students who were asked to predict the outcomes of basketball games were *more* confident when given semantically rich information (the names of the teams playing in the game) in addition to plain statistical cues (e.g. season records) than when not given such name information. The judges who had been given name information were in fact *less* accurate in their predictions. Future research should investigate further the commonalities across the informational setups that produce such dissociations.

Limitations and reservations

The blindfolding procedure is a useful experimental tool for investigating the process of combining opinions, yet it may be less applicable in realistic settings. Other procedures that enable decision makers to suspend their personal opinions may prove useful. Consider the nominal group method, in which no single member dominates the outcome of the process (Rowe & Wright, 2001). In a typical application of this method (Fraser, Pilpel, Kosecoff, & Brook, 1994), an expert panel of physicians was charged with creating guidelines for deciding on gallbladder surgery. To prevent undue influence, the group administrator was restricted from expressing an opinion and was simply supposed to integrate panel members' opinions by using an averaging formula. The administrator could be conceived as blindfolded.

Our studies are also limited in that they use only numerical judgments, yet we think that their conclusion can be extended to broader contexts, such as ones involving categorical judgments. A study by Greitemeyer and Schulz-Hardt (2003) found that participants' initial choice of one discrete alternative (out of three) reduced their ability to utilize the opinions of other group members and thereby improve their final choice. This study, illustrating the influences of egocentrism in making discrete choices, implies that decision makers could mitigate such egocentric biases by suspending their choices.

Finally, suspension of judgment may be less important in settings where the overweighting of one's own opinion would be warranted, such as when the decision maker's expertise is superior to that of her advisors. Yet, in numerous applied settings, individuals consult with peers and members of the same social and professional circles who tend to be similarly knowledgeable—for example, a group of students, a panel of doctors, or a committee of economic advisors (Armstrong, 2001; Schrah, Dalal, & Sniezek, 2006). In such groups, valid differences in expertise are hard to establish. Decision makers therefore need to have fuller appreciation of the detrimental effects of egocentric discounting. Devising methods for counteracting such self–other biases should allow people to make better use of the wisdom of others.

ACKNOWLEDGEMENTS

This research was supported by Grant Nos. 344/05 and 327/10 from the Israel Science Foundation to I. Yaniv and the Hebrew University Presidential Doctoral Fellowship to S. Choshen-Hillel. The assistance of Maxim Milyavsky is gratefully acknowledged.

DOI: 10.1002/bdm

REFERENCES

Armstrong, J. S. (2001). Combining forecasts. In J. S. Armstrong (Ed.), *Principles of forecasting: A handbook for researchers and practitioners*. Norwell, MA: Kluwer.

Bonaccio, S., & Dalal, R. S. (2006). Advice taking and decision making: An integrative literature review, and implications for the organizational science. *Organizational Behavior and Human Decision Processes*, *101*, 127–151.

Budescu, D. V., Rantilla, A. K., Yu, H. T., & Karelitz, T. M. (2003). The effects of asymmetry among advisors on the aggregation of their opinions. *Organizational Behavior and Human Decision Processes*, *90*, 178–194.

Budescu, D. V., & Yu, H.-T. (2007). Aggregation of opinions based on correlated cues and advisors. *Journal of Behavioral Decision Making*, *20*, 153–177.

Cohen, G. L., Aronson, J., & Steele, C. M. (2000). When beliefs yield to evidence: Reducing biased evaluation by affirming the self. *Personality and Social Psychology Bulletin*, *26*, 1151–1164.

Edwards, W. (1968). Conservatism in human information processing. In B. Kleinmuntz (Ed.), *Formal representations of human judgment* (pp. 17–52). New York: John Wiley & Sons.

Einhorn, H. J., Hogarth, R. M., & Klempner, E. (1977). Quality of group judgment. *Psychological Bulletin*, *84*, 158–172.

Fraser, G. M., Pilpel, D., Kosecoff, J., & Brook, R. H. (1994). Effect of panel composition on appropriateness ratings. *International Journal for Quality in Health Care*, *6*, 251–255.

Gino, F., Shang J., & Croson, R. (2009). The impact of information from similar or different advisors on judgment. *Organizational Behavior and Human Decision Processes*, *108*, 287–302.

Greitemeyer, T., & Schulz-Hardt, S. (2003). Preference-consistent evaluation of information in the hidden profile paradigm: Beyond group-level explanations for the dominance of shared information in group decisions. *Journal of Personality and Social Psychology*, *84*, 322–339.

Hall, C. C., Ariss, L., & Todorov, A. (2007). The illusion of knowledge: When more information reduces accuracy and increases confidence. *Organizational Behavior and Human Decision Processes*, *103*, 277–290.

Harvey, N., & Fischer, I. (1997). Taking advice: Accepting help, improving judgment and sharing responsibility. *Organizational Behavior and Human Decision Processes*, *70*, 117–133.

Harvey, N., & Harries, C. (2004). Effects of judges' forecasting on their later combination of forecasts for the same outcomes. *International Journal of Forecasting*, *20*, 391–409.

Herzog, S. M., & Hertwig, R. (2009). The wisdom of many in one mind: Improving individual judgments with dialectical bootstrapping. *Psychological Science*, *20*, 231–237.

Klayman, J., & Ha, Y-W. (1987). Confirmation, disconfirmation, and information in hypothesis testing. *Psychological Review*, *94*, 211–228.

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.

Mannes, A. E. (2009). Are we wise about the wisdom of crowds? The use of group judgments in belief revision. *Management Science*, *55*, 1267–1279.

Ronis, D. L., & Yates, J. F. (1987). Components of probability judgment accuracy: Individual consistency and effects of subject matter and assessment method. *Organizational Behavior and Human Decision Processes*, *40*, 193–218.

Rowe, G., & Wright, G. (2001). Expert opinions in forecasting: The role of the Delphi technique. In J. S. Armstrong (Ed.), *Principles of forecasting: A handbook for researchers and practitioners*. Norwell, MA: Kluwer Academic Publishers, 125–144.

Schrah, G. E., Dalal, R. S., & Sniezek J. A. (2006). No decision-maker is an island: Integrating expert advice with information acquisition. *Journal of Behavioral Decision Making*, *19*, 43–60.

Soll, J. B., & Larrick, R. P. (2009). Strategies for revising judgment: How (and how well) people use others' opinions. *Journal of Experimental Psychology. Learning, Memory, and Cognition*, *35*, 780–805.

Soll, J. B., & Mannes, A. E. (2011). Judgmental aggregation strategies depend on whether the self is involved. *International Journal of Forecasting*, *27*, 81–102. doi: 10.1016/j.ijforecast.2010.05.003

Strack, F., & Mussweiler, T. (1997). Explaining the enigmatic anchoring effect: Mechanisms of selective accessibility. *Journal of Personality and Social Psychology*, *73*, 437–446.

Surowiecki, J. (2004). *The wisdom of crowds*. New York: Doubleday.

Tetlock, P. E. (2005). *Expert political judgment*. Princeton, NJ: Princeton University Press.

Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, *185*, 1124–1131.

Vul, E., & Pashler, H. (2008). Measuring the crowd within: Probabilistic representations within individuals. *Psychological Science*, *19*, 645–647.

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., Choshen-Hillel, S., & Milyavsky, M. (2009). Spurious consensus and opinion revision: Why might people be more confident in their less accurate judgments? *Journal of Experimental Psychology. Learning, Memory, and Cognition*, *35*, 558–563.

Yaniv, I., Choshen-Hillel, S., & Milyavsky, M. (2011). Receiving advice on matters of taste: Similarity, majority influence, and taste discrimination. *Organizational Behavior and Human Decision Processes*, *15*, 111–120.

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.

Authors' biographies:

Ilan Yaniv received his PhD from the University of Michigan, taught at the University of Chicago and is now on the faculty of the Hebrew University.

Shoham Choshen-Hillel is a graduate student at the Hebrew University of Jerusalem. She is interested in advice-based decision making and in opinion revision processes, as well as in social preference, inequality aversion and prosocial behaviors.

Authors' address:

Ilan Yaniv and Shoham Choshen-Hillel, Department of Psychology & Center for the Study of Rationality, Hebrew University of Jerusalem, Jerusalem 91905, Israel.

DOI: 10.1002/bdm