Ilan Yaniv & Shoham Choshen-Hillel

Department of Psychology & Center for the Study of Rationality
Hebrew University of Jerusalem, Israel

Exploiting the wisdom of others: A bumpy road to better decision making

This paper is part of the Decision Making for a Social World webconference

Organized by the International Cognition and Culture Institute and
the Philosophy, Politics and Economics Program at the University of Pennsylvania

The paper is followed by a **discussion**

Introduction

It is 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 so as to improve one's final decision (Bonaccio & Dalal, 2006; 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 the others' advice. For example, the student (or manager) seeking others' opinions might or might not have the information to form a confident, independent preliminary forecast.

How might the presence or 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 & Yu, 2007). Normative studies have shown that statistical equal-weighting of judgments yields aggregate forecasts that are more accurate than the individual opinions upon which they are based (Larrick & Soll, 2006). 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; 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. While 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 (Mannes, 2009). The accuracy costs of this self-other bias are considerable. They become even greater when *several* other opinions are available, since participants assign the other opinions about the same weight altogether, 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

---

This presentation is an abbreviated version of Yaniv and Choshen-Hillel (in press). We thank Maxim Milyavsky for his contribution. 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. Please email correspondence to ilan.yaniv@huji.ac.il

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).

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, due to their biased weighting of evidence (Lord, Ross, & Lepper, 1979). In sum, the findings in the area of advice taking are consistent with the general conclusion that one's incorporation of new evidence is biased in the direction of one's prior beliefs.

Our study 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), while the ones assigned to the *blindfold* condition could not. We evaluated participants' policies for weighting other opinions and their judgmental accuracy. 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 participants who did *not* form prior opinions to gain more from the advice presented to them. While 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

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, in addition, 1 IS (about $0.25) for each estimate that fell within the range extending 15% on either side of the correct answer.

As shown in Table 1, there were two between-participants conditions (full-view vs blindfold). The *full-view* condition began with a practice part that included 5 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. 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.

Table 1: Sample materials and conditions in study.

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

*The first listed estimate was sampled in the blindfold condition.

The advisory estimates were sampled at random, with labels such as #26, #4, and #19, 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. 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*. On each trial, the participants were asked to estimate the caloric value of the (concealed) target food.

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), $t(33) = 7.49$, $p < .01$, thereby replicating the standard finding that receiving advice is beneficial. Participants' final estimates were even more accurate in the blindfold than in the full-view condition (66.0 vs 76.2), $t(65) = 2.55$, $p < .05$.

Do the full-view participants perform less well because they discount advice? The evidence of egocentric discounting of advice comes from two measures. First, full-view participants made final judgments that were *identical* to their initial ones in 52.5% of the cases, whereas blindfolded

participants adopted their first-listed estimates in only 12.3% of the cases, $t(65) = 9.10$, $p < .001$. 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), $t(65) = 14.16$, $p < .001$. These results provide evidence of egocentric discounting in the full-view condition.

Table 2: Results of Study.

	Full-view (n = 34)	Blindfold (n = 33)
Accuracy (mean absolute errors)		
Initial estimate	91.2	--
Final estimate	76.2	66.0
Average (equal weight) of 6 estimates	55.2	54.9
Measures of egocentrism		
% keeping the first-listed estimate*	52.5	12.3
Distance between final and first-listed estimate*	26.6	91.3
Distance between final estimate and average of 6 estimates	53.8	37.9

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

Subsequent analyses further corroborate the idea that egocentric bias is detrimental to accuracy. Averaging the 6 opinions produced more accurate estimates than did either the full-view judges (55.2 vs 76.2), $t(33) = 6.90$, $p < .001$, or the blindfolded judges (54.9 vs 66.0), $t(32) = 6.10$, $p < .001$. 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), $t(65) = 5.50$, $p < .001$. The greater proximity of the blindfolded judges to equal weighting presumably accounts for their relatively greater accuracy.

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.

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. Two patterns of findings are particularly informative here – first, the effects of blindfolding on accuracy; second, the judges' policies for using advice.

Blindfolding effects on accuracy. Our participants were asked to make decisions based on other people's advisory opinions. In one condition, participants could form prior estimates (full-view condition), while in the other, they could not form their own personal opinions, but had to rely only on advice (blindfold condition). The blindfolded judges used the advice better and gained more in accuracy than did the full-view judges. 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 estimates enhances their use of the advice and their final accuracy.

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, since 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, since 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, while in the remaining cases they made final estimates that were close to their initial estimates. 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.

Why might judges persist in their opinions? Why did the full-view judges persist in their initial estimates rather than take full advantage of the other opinions? We suggest the following account, based on the decision makers' differential awareness of people's reasons for making choices. In the process of forming their final opinions, decision makers 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. The full-view participants formed an explicit opinion and assigned more weight to it than to alternative advisory opinions. Thus suspending judges' ability to form personal opinions until after they had produced blindfolded (i.e., advice-based) estimates leads them to more egalitarian use of the advice.

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 under-weighted the informational value of the samples obtained in the task. Edwards (1968, p.17) 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." Our own documentation of informational waste is consistent with Edwards's 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.

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, an expert panel of physicians was charged with creating guidelines for deciding on a specific surgery. To prevent undue influence, the group administrator was restricted from expressing an opinion and was simply supposed to integrate panel members' opinions 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 conclusions can be extended to 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.

In sum, decision makers need to have fuller appreciation of the detrimental effects of egocentric discounting. Devising methods for counteracting such self-other biases, such as suspending judgment, should allow people to make better use of the wisdom of others.

References

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

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.

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.

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

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 Armstrong, J. S. (Ed.), *Principles of forecasting: A handbook for researchers and practitioners*. Norwell, MA: Kluwer Academic Publishers, 125-144.

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.

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.

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, 115,* 111-120.

Yaniv, I., & Choshen-Hillel, S. (in press). Exploiting the wisdom of others to make better decisions: Suspending judgment reduces egocentrism and increases accuracy. *Journal of Behavioral Decision Making*.