

On the Evaluation of One-sided Evidence

LYLE A. BRENNER, DEREK J. KOEHLER AND AMOS TVERSKY
Stanford University, USA

ABSTRACT

We examine predictions and judgments of confidence based on one-sided evidence. Some subjects saw arguments for only one side of a legal dispute while other subjects (called 'jurors') saw arguments for both sides. Subjects predicted the number of jurors who favored the plaintiff in each case. Subjects who saw only one side made predictions that were biased in favor of that side. Furthermore, they were more confident but generally less accurate than subjects who saw both sides. The results indicate that people do not compensate sufficiently for missing information even when it is painfully obvious that the information available to them is incomplete. A simple manipulation that required subjects to evaluate the relative strength of the opponent's side greatly reduced the tendency to underweigh missing evidence.

KEY WORDS: confidence; uncertainty; bias; incomplete information

In daily life we are often exposed to one-sided information about various issues, disputes, and controversies involving two or more parties. Sometimes we encounter one-sided information and incorrectly believe that it is complete or fully representative of the total body of evidence. However, there are many cases where one is faced with one-sided information knowing full well that it represents only one side of a dispute. In such cases it is obvious that one should compensate for the position of the other side, but it is less obvious how to do so.

This problem often arises in legal disputes, where each side does not know the details of the specific evidence and arguments to be presented by the other side. Observers of mock trials have noted that clients are often surprised and dismayed by the strength of the opposition's case as simulated in the mock trial, even when they know in advance the general form of the evidence to be presented by the opposition. Evidently, these clients have underestimated the impact of the opponent's position, and have consequently overestimated their chances of winning the case. In order to predict accurately the outcome of a legal process, judgments of the strength of one's own side must be tempered by an evaluation of the other side's strength. This problem arises in other contexts as well. Forming an opinion about an unfamiliar political conflict (e.g. between Yemen and South Yemen) or about a messy divorce on the basis of an account presented by one of the disputants illustrates the problem of evaluating an issue on the basis of one-sided evidence.

The present studies investigate people's ability to deal with incomplete information in situations in which it is obvious that the available evidence is one-sided, there is no personal involvement in either side, and there is no reason to believe that the arguments presented by one side are more credible than

those presented by the other side. The experiments below involve prediction of the votes of a jury for a series of legal cases. Subjects received either background information only, background and plaintiff information, background and defendant information, or full (background, plaintiff, and defendant) information. Subjects then predicted the number of fully informed jurors who voted for a given side, and also provided a measure of confidence in their predictions.

In order to facilitate the task of predicting the jury vote on the basis of one-sided evidence, we constructed a task with the following characteristics:

- (1) Subjects who received one-sided information were explicitly told (repeatedly, both in written and verbal instructions) that the jury to be predicted received the arguments for both sides.
- (2) Subjects who received one-sided information were told that the participants in the study were randomly assigned to one side or another (or to the jury). Hence, the information that they received had no special status.
- (3) The arguments for each of the sides did not include any new facts about the case beyond what was provided in the background information.
- (4) Furthermore, the arguments for both sides were fairly straightforward and quite predictable from the background information. No surprising deductions or ingenious re-interpretations of the facts were present in the arguments for either side.
- (5) Subjects were asked to predict a concrete, measurable outcome: namely, the number of subject-jurors finding in favor of one side.

The question arises: how should subjects predict on the basis of one-sided evidence? Obviously, they should keep in mind the discrepancy between the information available to them and the information available to the jury whose judgments they are predicting. Two cognitive strategies could be employed to deal with this discrepancy. First, subjects might try to construct arguments that are likely to be used by the opponent. Consider, for example, a child custody dispute. If one receives the father's argument that he has greater financial resources and thus can better provide for the child, one could imagine the possible response by the mother (e.g. that she can devote more time to the child and could give the child more attention). Based on the argument at hand and the other in mind, one could form an estimate of the jury's likely response to the case. Alternatively, one could try to assess the relative strength of the opposing position (e.g. that the mother would seem to have the upper hand), without considering in detail the specific arguments for either side. Naturally, a combination of the two strategies could also be employed. Thus, the generation of specific arguments for the other side or the assessment of the relative strength of the two sides could be used to make predictions based on one-sided evidence.

Subjects who predict the jury vote given one-sided evidence should (normatively) be less confident in their predictions than subjects who receive all the information available to the jury. This follows from the fact that in addition to the uncertainty regarding the judgment of the jury, subjects who receive one-sided evidence face additional uncertainty regarding the arguments of one of the two sides. The additional uncertainty will generally reduce the accuracy of predictions, and should therefore reduce the subject's confidence in these predictions.

In contrast, it has been proposed that people's confidence in their prediction generally increases with the consistency or coherence of the available information (Kahneman and Tversky, 1973; Peterson and Pitz, 1988). In the present study each side presents a coherent position that is inconsistent with the position of the other side. Consequently, subjects who receive consistent one-sided evidence may be more confident in their predictions than subjects receiving conflicting information from both sides. The following studies investigate the effects of one-sided evidence on the bias, accuracy, and confidence associated with the predictions of the jury vote.

EXPERIMENT 1

Method

Subjects

Participants were 137 students at San José State University who filled out questionnaires in small groups of five to ten individuals. They received course credit for their participation.

Materials

Six legal scenarios were used. The scenarios were based on actual legal cases which were simplified and changed as appropriate. For each case there were three sets of information: background information, plaintiff arguments, and defendant arguments. Each set of information was one or two paragraphs long. The six cases consisted of two civil cases, two criminal cases, and two child custody cases. For all subjects, cases of the same type (e.g. civil, criminal, custody) were presented consecutively. The order of presentation of the three case types was counterbalanced.

For convenience, we use 'plaintiff' to refer to the plaintiff for civil cases, the prosecution for criminal cases, and the mother for custody cases. Similarly, 'defendant' refers to the defendant for civil and criminal cases, and the father for custody cases. The arguments for the plaintiff and the defendant were written so as to stand alone, with no reference to the arguments of the other side. The background material consisted of one-paragraph summaries of the central issues of the cases, with no arguments for either side. No new facts or evidence beyond those stated in the background information were presented in the plaintiff's or defendant's arguments. The background information and arguments for a sample case are presented in the appendix.

Procedure

Subjects were randomly assigned to one of two groups: jury ($n = 34$) or non-jury ($n = 106$). Non-jury subjects had partial (i.e. one-sided or background-only) information for all six cases. Jury subjects received full (i.e. two-sided) information and actually voted for one of the two sides in each case. Subjects made predictions for all six cases.

For the non-jury subjects, the instructions read:

On the following pages you will find information about six different court cases. A previous group of subjects was given three sets of information about each case: They were given some background information describing the general circumstances of the case, and they were given a summary of the arguments presented by each side. Based on this information, each subject gave his or her judgment on the case by ruling in favor of one of the two sides.

Of these subjects, we have selected a group of twenty at random. Your task is to estimate how many of the twenty favored a given side for each case. Because we are interested in judgments made based on partial information, however, you will only be given a subset of the information presented to our original subjects. For some of the cases you will be given the background information and the arguments made by one side, but not the arguments made by the other side. For other cases you will only be given the background information, without the arguments presented by either side. In all cases, your task will be to give your best guess as to how many of the twenty subjects ruled in favor of a given side.

The subjects were told that the jury they were predicting did not deliberate or make a decision as a group; the jury was defined as a group of 20 other subjects who individually voted on the cases after reading background information and arguments for both sides. All subjects were asked to indicate

their *best guess* as to how many of the 20 jury members voted for the plaintiff by circling a number between 0 and 20.

Subjects were also asked to set an *uncertainty range* (similar to a confidence interval) around their best guess. They were instructed to make high and low estimates such that they felt 90% confident that the actual number of plaintiff-voting jurors fell inside the range between the high and low estimates. More specifically, they were told that the actual jury vote should fall below their low estimate 5% of the time and above their high estimate 5% of the time. The difference between the subject's high and low estimates (i.e. the size of the subject's uncertainty range) is interpreted as a measure of confidence in the best guess, with narrower ranges indicating more confidence and wider ranges indicating less confidence. For example, a subject with a low estimate of 5 and a high estimate of 15, yielding an uncertainty range of 10, is seen as less confident in his or her prediction than a subject with a low estimate of 8 and a high estimate of 12, yielding an uncertainty range of 4.

Jury subjects received essentially the same instructions, except that no mention was made of partial information. For reasons of economy, jury subjects both voted *and* predicted. They were asked to predict the votes of twenty other jurors who, like them, had read arguments for both sides, and to form a 90% uncertainty range around their best guesses. Before making their predictions, they voted for one side in each case. (Using the same subjects to perform both tasks is justified by the results of Experiment 2 discussed below.)

For the non-jury subjects, the first two cases always consisted of background information only. Following these cases, approximately half the subjects made predictions for two Plaintiff-only cases, and then for two Defendant-only cases, and the other half received two Defendant-only cases followed by two Plaintiff-only cases. Non-jury subjects were reminded for each case that the jury being predicted had access to more information than they did.

Jury subjects read background, plaintiff and defendant arguments for all six cases. Half of the jury subjects read plaintiff arguments first and the other half read defendant arguments first. Because this variable had no reliable effect, it is omitted from further analysis.

Results and discussion

Because of the complexity of the design, two separate repeated-measures analysis of variance (ANOVA) models were fit, one for the non-jury subjects and one for the jury subjects. Comparisons between jury and non-jury data are made across the ANOVA models using approximate z-tests based on the variance estimates obtained from the ANOVA models.

There are two dependent measures of interest: best guess and uncertainty range. Best guesses are coded in terms of the estimated number of jury members finding for the *plaintiff*. Thus, higher estimates indicate belief in a more plaintiff-prone jury, whereas lower estimates indicate belief in a defendant-prone jury. The hypothesis of insufficient adjustment for the unavailable side would be supported by a pattern of higher estimates for the Plaintiff-only condition and lower estimates for the Defendant-only condition.

We can define three measures of bias of the best guesses: Plaintiff-bias, Defendant-bias, and Total-bias. Plaintiff-bias is the extent to which predictions of the number of plaintiff-voting jurors under Plaintiff-only information exceed the predictions under full information (i.e. predictions of the Jury). Defendant-bias, similarly, is the extent to which predictions of the number of plaintiff-voting jurors under Defendant-only information fall below the jury predictions. Total-bias is the sum of these two bias measures, or more simply the difference between mean Plaintiff-only and mean Defendant-only best guesses.

Mean best guesses for the four information conditions and the six cases are displayed in Exhibit 1, and the overall mean best guesses collapsing across cases are displayed in Exhibit 2. Mean estimates

Exhibit 1. Mean estimates of number of jury votes (out of 20) in favor of plaintiff, by information condition and case for Experiment 1

Case	Background	Information condition		Jury
		Plaintiff-only	Defendant-only	
1	11.8	15.1	9.2	12.1
2	10.3	10.7	9.4	10.5
3	11.3	11.7	10.8	11.1
4	8.3	11.4	10.9	12.0
5	11.8	12.8	8.6	11.7
6	8.6	11.6	7.9	10.6

Exhibit 2. Means (and standard deviations), by condition, for best guess of number of 20 jury members favoring plaintiff, uncertainty range (difference between high and low estimates of jury vote), and signed and absolute error between predicted and actual proportion of jury votes for plaintiff, for Experiment 1

Measure	Background	Information condition		Jury
		Plaintiff-only	Defendant-only	
Best Guess	10.3 (4.2)	12.3 (3.9)	9.5 (4.0)	11.3 (3.5)
Uncertainty Range	10.8 (2.5)	10.5 (3.0)	10.3 (2.9)	11.1 (3.3)
Signed Error	0.03 (0.25)	0.13 (0.23)	-0.01 (0.22)	0.09 (0.20)
Absolute Error	0.21 (0.14)	0.22 (0.15)	0.18 (0.13)	0.18 (0.12)

for the Plaintiff-only condition are greater than the mean estimates for the Defendant-only condition for all six cases, and the overall contrast comparing the two conditions is highly significant, $t(659) = 7.9, p < 0.001$. Mean best guesses for the Plaintiff-only condition are greater than those for the Jury, $z = 2.76, p < 0.01$, and mean best guesses for the Defendant-only condition are less than those for the Jury, $z = 5.56, p < 0.01$. Background and Jury estimates fall in between the Plaintiff-only and Defendant-only estimates for five of the six cases. Thus, subjects who received one-sided evidence were biased in the direction of this evidence.

Mean uncertainty ranges for the four information conditions are also shown in Exhibit 2. Collapsing across the six cases, we find that uncertainty ranges for the Jury condition are significantly larger ($M = 11.1$) than uncertainty ranges for the one-sided information conditions ($M = 10.4$, $z = 3.61, p < 0.001$). Thus, full information produced wider uncertainty ranges than partial information. This observation indicates that people expected to be more accurate (i.e. exhibited greater confidence in their best guess) in the one-sided conditions rather than the jury condition.

Because each participant made judgments for six different cases, we can test for effects of the order in which the cases were judged. There were no effects of case order on best guesses for either the non-jury or jury groups, $F_s < 1$. However, uncertainty range widths increased with case order for the jury group ($F(1,144) = 5.42, p < 0.05$), and decreased with case order for the non-jury group ($F(1,493) = 7.0, p < 0.01$), when adjusting for the different information conditions encountered in later cases. Jury participants became less confident in their predictions, and non-jury participants more confident, as they encountered new cases.

To assess the accuracy of predictions, the actual proportion of jurors voting for the plaintiff was compared to the predicted proportion in each case. The mean difference between the best guess proportion and the target proportion is presented in Exhibit 2 as 'Signed Error'. The mean of the

absolute value of the difference between predicted and actual jury vote is also presented in Exhibit 2 as 'Absolute Error'. We use this latter measure as an index of predictive accuracy.

Mean Jury accuracy was reliably better than mean Plaintiff-only accuracy ($z = 3.75$, $p < 0.01$), whereas mean Defendant-only accuracy matched mean Jury accuracy ($z = 0.09$). In summary, full information yielded equally accurate or more accurate predictions than did partial information, but was accompanied by wider uncertainty ranges.

The results show that our subjects did not sufficiently adjust for the unavailable information. However, the discrepancy between mean Plaintiff-only and mean Defendant-only estimates varied considerably across the six cases. Furthermore, there was a substantial correlation ($r = -0.40$) between Total-bias (mean Plaintiff-only best estimates minus mean Defendant-only best estimates) and the size of the Jury's majority vote. When the jury is evenly split (i.e. the majority is in the range 50% to 60%), the biasing effect of one-sided information is quite large. In contrast, when the jury is nearly unanimous, the effect is small. Ratings of the persuasiveness of the arguments for the two opposing sides of each case (obtained from subjects in Experiment 2) allow an independent assessment of the 'closeness' of the two sides. The degree of Total-bias correlated highly with the difference in rated argument persuasiveness between the two sides of a case ($r = -0.75$). Furthermore, there is a substantial correlation ($r = 0.46$) between Total-bias and the average persuasiveness of the two sides. Thus, the biasing effect of one-sided evidence is most pronounced in high-conflict situations, where there are strong and equally persuasive arguments for both sides.

EXPERIMENT 2

In the preceding experiment, subjects in the Jury group acted as predictors and as targets, both voting for one of the parties and judging how many other jurors would vote for the plaintiff. Experiment 2 was designed to test whether voting affects the jury's subsequent predictions, and, especially, the width of their uncertainty ranges.

Method

Subjects

Subjects were 54 students in an introductory psychology class at Stanford University who completed the task for course credit.

Procedure

Subjects received either a Voting packet or a Nonvoting packet by random assignment. Both packets contained the six legal cases used in Experiment 1, with the arguments for both sides presented for each case. After reading the arguments for a case, Voting subjects voted for one side, and then made predictions about the number of plaintiff-voting jurors. This is the same sequence of tasks performed by the jury subjects in Experiment 1. Nonvoting subjects made their predictions without first voting.

After making their predictions, all subjects rated the persuasiveness of each side's arguments on a 9-point scale with endpoints labeled 'not at all persuasive' and 'very persuasive'.

Results

Voting and Nonvoting subjects did not differ systematically in terms of their best guesses; the Voting group made higher mean estimates for three of the cases, and lower estimates for the other three cases.

Collapsing across cases, the Voting jurors were slightly more plaintiff-oriented ($M = 11.4$) than the nonvoting jurors ($M = 10.6$), but the difference was not statistically significant ($F(1,52) = 3.3$, $p > 0.05$). Mean uncertainty ranges were identical ($M = 8.8$) for both groups. This result justifies the use of the same subjects for both voting and predicting.

EXPERIMENT 3

The preceding experiments differ from many real-life encounters in that the information and the arguments were not very extensive, and they were presented in written form. It is of interest, therefore, to test whether the bias induced by one-sided evidence is present when people are exposed to more extensive argumentation, presented orally by the disputants. In this experiment, subjects viewed videotape of 'trial lawyers' delivering fairly extensive arguments for the two sides.

Method

Subjects

Participants were 203 students from San José State University who completed the experiment in groups of 10–30 people, for course credit.

Materials

We videotaped two actors presenting expanded arguments for the two civil cases used in Experiment 1 (cases 1 and 2). Background statements for each case were read by a third actor. One actor presented the Plaintiff arguments for both cases, and the other presented the Defendant arguments for both cases. Each subject saw each 'lawyer' exactly once.

The background statement lasted approximately one minute, while the arguments for the two sides lasted nearly four minutes each. The actors stood behind a podium and delivered their arguments to the camera, in a manner resembling a lawyer's closing arguments.

Procedure

Subjects were told that they were to view arguments based on legal cases and make predictions about how other subjects voted on the cases. Subjects saw information for the two cases, in one of four sequences: Background/Jury, Jury/Background, Plaintiff-only/Defendant-only or Defendant-only/Plaintiff-only. Thus, half the subjects saw one-sided arguments for both cases, and half saw full information for one case and only background information for the other. The cases were presented in the same order for all subjects, preventing the analysis of order effects.

Subjects had a written transcript of the background statement so that they could follow along. No transcripts of the arguments for the two sides were provided to the subjects, in order to induce the subject to attend to the videotape.

As in Experiment 1, subjects who received one-sided evidence were explicitly told that they were receiving partial information (both in written instructions and by the experimenter), and were asked to predict the number of jury members who voted for the plaintiff. Subjects recorded their predictions and uncertainty ranges after viewing the appropriate arguments for each case.

Exhibit 3. Mean estimates of number of jury votes (out of 20) in favor of plaintiff, by information condition and case for Experiment 3

Case	Background	Information condition		Jury
		Plaintiff-only	Defendant-only	
1	13.0	13.8	7.6	9.6
2	9.6	11.5	7.9	10.2

Results

The data were analyzed using a repeated measures ANOVA model with the sequence of information conditions (four levels) as a between-subjects factor and case (two levels) as a within-subjects factor.

Predictions were once again biased in the direction of one-sided information; mean estimates of the jury vote are displayed in Exhibit 3. Overall, the mean Plaintiff-only best guess ($M = 12.6$) exceeded the mean Jury best guess ($M = 9.9$, $F(1,198) = 21.5$, $p < 0.001$) which exceeded the mean Defendant-only best guess ($M = 7.7$, $F(1,198) = 14.4$, $p < 0.001$). The size of the Total-bias was substantially larger than in Experiment 1, nearly 5 points on a 21-point scale.

Uncertainty ranges were again wider for Jury subjects ($M = 8.6$) than for subjects in one-sided conditions ($M = 7.9$, $F(1,198) = 4.8$, $p < 0.05$), and the size of the effect is nearly identical to that found in Experiment 1. Jury estimates were again more accurate by the absolute error measure ($M = 0.22$) than were those for the one-sided conditions ($M = 0.24$), although this difference was not statistically significant.

EXPERIMENT 4

In order to make reasonable predictions on the basis of one-sided evidence one must in some way take into account the position of the other side. As we have suggested earlier, this could be done either by generating specific arguments that are likely to be presented by the other side, or by assessing the overall relative strength of the two sides. The results of the preceding studies suggest that subjects did not effectively employ either of these strategies, and hence the application of these strategies could lead to less biased predictions. In Experiment 4 we investigate the effect of an extremely simple debiasing procedure (cf. Fischhoff, 1982), which merely asks subjects to rate the relative strength of the two sides prior to the prediction of the juror vote.

Method

Subjects

Participants were 149 students in an introductory psychology class at Stanford University who completed the task for course credit.

Procedure

Each subject received a packet containing instructions and the two civil case scenarios used in Experiments 1 and 3. All subjects received one-sided arguments only, Plaintiff-only for one case and Defendant-only for the other case.

Half of the subjects (the Standard condition) completed the usual task of predicting the vote of a jury that saw both sides. The other half of the subjects (the Evaluation condition) completed this prediction task after evaluating the relative strength of the other side. The instructions read:

Now you've read the arguments for the plaintiff. Recall that the 'jury' subjects read both the plaintiff's arguments and the defendant's arguments. On the basis of what you've read above, do you expect the defendant's arguments to be weaker or stronger than the plaintiff's arguments?

The instructions for the Defendant-only condition were parallel. Subjects rated the relative strength of the other side's arguments by marking one of five boxes labeled 'much weaker', 'slightly weaker', 'equal', 'slightly stronger', and 'much stronger'.

After answering this question and estimating the number of jurors finding for the plaintiff, the Evaluation subjects answered the following question, marking one of five boxes:

Would you expect to be more or less confident in your estimate if you had seen the arguments for both sides?

Results

Mean estimates of jury vote for the Standard and Evaluation conditions are displayed in Exhibit 4. Hypothesis tests are based on a repeated measures ANOVA model including Evaluation/Standard (two levels) and sequence of information conditions (two levels) as between-subjects factors, and case (two levels) as a within-subjects factor.

Exhibit 4. Mean estimates of number of jury votes (out of 20) in favor of plaintiff, by condition (Evaluation or Standard), case, and side seen (Plaintiff or Defendant) for Experiment 4. In the Evaluation condition, participants rated the relative strength of the other side before predicting the jury vote

Condition	Case	Side seen		Difference
		Plaintiff	Defendant	
Evaluation	1	13.9	10.1	3.8
	2	8.6	6.8	1.8
Standard	1	15.7	5.9	9.8
	2	10.5	5.6	4.9

Overall, the mean Plaintiff-only estimate is greater than the mean Defendant-only estimate, $F(1,81) = 68.1, p < 0.001$, and the size of this difference varies across the Standard and Evaluation conditions, $F(1,81) = 13.3, p < 0.001$. The difference (averaged across cases) is much smaller (2.7) when subjects evaluate the relative strength of the two sides compared to when they do not (7.1). Thus, simply rating the expected relative strength of the other side serves to markedly reduce the biasing effect of one-sided information.

It is worth noting that the unavailable other side was, on average, rated as only slightly weaker than the available side (mean = 2.8 on a 5-point scale where 3 indicates equal strength). This observation suggests that the bias is driven primarily by a failure to carefully consider the other side, rather than by a tendency to underestimate its relative strength.

A majority of the subjects (60%) indicated that they would be more confident if they had seen both sides, while 15% expected to be equally confident, and 25% expected to be less confident. This pattern

indicates that most subjects had the correct normative intuition even though the preceding experiments show that uncertainty ranges tend to be tighter with partial rather than complete information.

GENERAL DISCUSSION

The results of the present experiments may be summarized as follows. First, subjects did not adjust sufficiently for one-sided evidence. Second, predictions based on one-sided information were associated with tighter uncertainty ranges (which indicate greater confidence in the accuracy of the prediction) despite the fact that their accuracy was generally lower. Third, the biasing effect of one-sided evidence is most pronounced when the two sides are evenly balanced and both sides are judged to be highly and equally persuasive. Finally, the biasing effect can be greatly reduced by encouraging an evaluation of the relative strength of the two sides.

Our findings are consistent with those of studies documenting insensitivity to sample bias in judgment (e.g. Nisbett and Borgida, 1975; Hamill, Wilson, and Nisbett, 1980; Ross, Amabile, and Steinmetz, 1977; Nisbett and Ross, 1980). What distinguishes the present study, aside from the legal context, is the deliberate effort to encourage subjects to consider the missing information, and the impact of missing information on uncertainty ranges. As noted earlier, incomplete information should (normatively) yield less confidence than full information in this task. Our subjects did not exhibit this pattern, and produced tighter uncertainty ranges when only one side was presented. Not surprisingly, Jury subjects, who had access to full information, were equally accurate or more accurate in their predictions. Thus, there was a negative correlation between confidence (as measured by the width of the uncertainty range) and accuracy (as measured by an absolute error metric). A similar relation between accuracy and confidence has been observed by Koehler *et al.* (1995) and by Peterson and Pitz (1986, 1988). The latter authors distinguished between the width of the uncertainty range, which we have interpreted as an index of confidence, and direct ratings of confidence in one's best prediction.

Our findings may seem at variance with the results of Oskamp (1965), who found that confidence in predictions increased with the amount of information available to subjects, while accuracy remained essentially unchanged. Oskamp's study, however, differed from ours in that subjects received additional information which was generally consistent with the previously presented information. In the present studies, in contrast, the arguments provided by one side conflict with the arguments provided by the other side. Similarly, Peterson and Pitz (1986, 1988) found that adding information about the performance of baseball teams that conflicted with previously available information (e.g. good pitching and poor hitting) generally increased the width of uncertainty ranges regarding the prediction of the number of season victories of these teams. Thus, people's confidence in the accuracy of their predictions depends primarily on the consistency or coherence of available information, which could give rise to a negative correlation between confidence and accuracy (Kahneman and Tversky, 1973).

Several studies of the evaluation of multi-attribute objects (e.g. Yates, Jagacinski, and Faber, 1978; Jagacinski, 1991; Jaccard and Wood, 1988) have suggested that missing information may be implicitly replaced with a below-average value on the relevant dimension. Thus, incompletely described objects are devalued relative to fully described objects. The present results are consistent with these findings in the sense that the impact of the missing side was underweighted in predictions of the jury vote.

More generally, there is evidence that information that is not made explicit is often neglected even when it can be readily retrieved or generated (e.g. Fischhoff, Slovic, and Lichtenstein, 1978; Johnson *et al.*, 1993; Tversky and Koehler, 1994). Perhaps the most striking feature of the present results is that the mere invitation to assess the relative strength of the two sides was sufficient to reduce the bias by more than a half. Instructions to generate the specific arguments for the other side, similar to

manipulations performed by Koriat, Lichtenstein, and Fischhoff (1980) and by Hoch (1985), may further reduce the bias induced by one-sided evidence.

APPENDIX

David Thornton v. Thrifty Drug Corp.

Background

On 3 September plaintiff David Thornton, a 43-year-old union field representative, was present in Thrifty Drug Store #168, performing a routine union visit. Within ten minutes of his arrival, a store manager confronted him and told him he could no longer speak with the union employees on the floor of the store. Instead, he would have to see them in a back room while they were on break. Such a request is allowed by the union contract with Thrifty Drug, but had never before been enforced. When Mr Thornton objected, he was told that he had the choice of conforming to these requirements, leaving the store, or being arrested. At this point, Mr Thornton indicated to the manager that he had always been allowed to speak to employees on the floor for as much as ten minutes, as long as no business was disrupted, and that he would rather be arrested than change the procedure of his routine visit. The manager then called the police and had Mr Thornton handcuffed in the store for trespassing. After he was booked and put into a holding cell for a brief time, all charges were dropped. Mr Thornton is suing Thrifty Drug for false arrest.

Plaintiff's arguments

Plaintiff Mr Thornton contends that the store manager had changed the rule under which he had been allowed, for three years prior to the event in question, to speak with employees while they were on the floor. He further argues that the defendant Thrifty Drug had no intention of prosecuting him at the time of arrest, and only intended to intimidate him from further vigorous representation of the union employees in the store. Using these legal maneuvers as a form of harassment constitutes false arrest. He claims that as a result of the false arrest, he was subjected to humiliation, anxiety, and emotional distress.

Defendant's arguments

Defendant Thrifty Drug Corporation contends that, according to Article 10 of the union contract, the store manager had the right to require that business agents meet with employees during their breaks, and off the store floor. Therefore, his demand that Mr Thornton conform to Article 10, leave, or be arrested was reasonable, and is allowed by the law. They further argue that their earlier policy of allowing interviews on the store floor had been suspended because of the amount of time these interviews took, and that the fact that the policy had been changed recently is irrelevant to the case at hand.

ACKNOWLEDGEMENT

This research was supported in part by grant MH-53046 from the National Institute of Health. We are grateful to Mark Forehand and Hae-June Ahn for research assistance.

REFERENCES

- Fischhoff, B. 'Debiasing', in Kahneman, D., Slovic, P. and Tversky, A. (eds), *Judgment under Uncertainty: Heuristics and biases*, Cambridge: Cambridge University Press, 1982.
- Fischhoff, B., Slovic, P. and Lichtenstein, S. 'Fault trees: Sensitivity of estimated failure probabilities to problem representation', *Journal of Experimental Psychology: Human Perception and Performance*, **4** (1978), 330–344.
- Hamill, R., Wilson, T. D. and Nisbett, R. E. 'Insensitivity to sample bias: Generalizing from atypical cases', *Journal of Personality and Social Psychology*, **39** (1980), 578–589.
- Hoch, S. J. 'Counterfactual reasoning and accuracy in predicting personal events', *Journal of Experimental Psychology: Learning, Memory and Cognition*, **11** (1985), 719–731.
- Jaccard, J. and Wood, G. 'The effects of incomplete information on the formation of attitudes toward behavioral alternatives', *Journal of Personality and Social Psychology*, **54** (1988) 580–591.
- Jagacinski, C. M. 'Personnel decision making: The impact of missing information', *Journal of Applied Psychology*, **76** (1991), 19–30.
- Johnson, E. J., Hershey, J., Meszaros, J. and Kunreuther, H. 'Framing, probability distortions, and insurance decisions', *Journal of Risk and Uncertainty*, **7** (1993), 35–51.
- Kahneman, D. and Tversky, A. 'On the psychology of prediction', *Psychological Review*, **80** (1973), 237–251.
- Koehler, D. J., Brenner, L. A., Liberman, V. and Tversky, A. 'Confidence and accuracy in trait inference: Judgment by similarity', *Acta Psychologica*, forthcoming (1995).
- Koriat, A., Lichtenstein, S. and Fischhoff, B. 'Reasons for confidence', *Journal of Experimental Psychology: Human Learning and Memory*, **6** (1980), 107–118.
- Nisbett, R. E. and Borgida, E. 'Attribution and the psychology of prediction', *Journal of Personality and Social Psychology*, **32** (1975), 932–943.
- Nisbett, R. E. and Ross, L. *Human Inference: Strategies and shortcomings of social judgment*, Englewood Cliffs, NJ: Prentice Hall, 1980.
- Oskamp, S. 'Overconfidence in case-study judgments', *The Journal of Consulting Psychology*, **29** (1965), 261–265.
- Peterson, D. K. and Pitz, G. F. 'Effects of amount of information on predictions of uncertain quantities', *Acta Psychologica*, **61** (1986), 229–241.
- Peterson, D. K. and Pitz, G. F. 'Confidence, uncertainty and the use of information', *Journal of Experimental Psychology: Learning, Memory and Cognition*, **14** (1988), 85–92.
- Ross, L., Amabile, T. M. and Steinmetz, J. L. 'Social roles, social control, and biases in social-perception processes', *Journal of Personality and Social Psychology*, **5** (1977), 485–494.
- Tversky, A. and Kahneman, D. 'Judgment under uncertainty: Heuristics and biases', *Science*, **185** (1974), 1124–1131.
- Tversky, A. and Koehler, D. J. 'Support theory: A nonextensional representation of subjective probability', *Psychological Review*, **101** (1994), 547–567.
- Yates, J. F., Jagacinski, C. M. and Faber, M. D. 'Evaluation of partially described multiattribute options', *Organizational Behavior and Human Performance*, **21** (1978), 240–251.

Authors' biographies:

Lyle Brenner received his PhD in Psychology from Stanford University in 1995. He is currently an associate at Decision Focus Incorporated in Mountain View, CA.

Derek Koehler received his PhD in Psychology from Stanford University in 1993, and is now assistant professor in the Department of Psychology at the University of Waterloo. His research interests include probability judgment and social cognition.

Amos Tversky received his PhD in Psychology from the University of Michigan in 1965, and now holds the Davis Brack Chair of Behavioral Sciences at Stanford University. His research interests include decision making and judgment under uncertainty.

Authors' address:

Lyle Brenner, Derek Koehler and Amos Tversky, Department of Psychology, Stanford University, Stanford, CA 94305-2130, USA.