Journal of Experimental Psychology:
Human Perception and Performance
1977, Vol. 3, No. 2, 349-358

Perceived Informativeness of Facts

Baruch Fischhoff
Decision Research, Eugene, Oregon

There are many tasks in which people are called on to disregard information that they have already processed. Dealing with inadmissible evidence in a courtroom setting, second-guessing the past, and responding to experimental psychologists' debriefing instructions are three tasks of this type; in all these cases, people have been found to experience considerable difficulty. The present experiment investigates these difficulties in a general form, using almanac-type questions. Subjects told the correct answers to such questions were found to overestimate both how much they would have known about the answer had they not been told and how much they actually did know about the answer before being told. Attempts to undo this *knew-it-all-along* effect by exhorting subjects to work harder or telling them about the bias failed. These results were discussed in terms of how the structure of one's knowledge is altered to accommodate new information.

Recent studies of hindsight (Fischhoff, 1975a, 1975b; Fischhoff & Beyth, 1975) have shown (a) that telling people that an event has occurred increases their subjective probability that it was going to happen and (b) that people underestimate the effect that hearing such reports has on their perceptions. Thus, they believe that they knew all along that the reported event was going to happen, even without the benefit of the report. This bias constitutes an underestimation of how surprising the report of what happened was and what one has to learn from it. If you "knew all along" what would happen, then you didn't need the report.

Underestimating what one learns from reported facts about past events may be a special case of underestimating what one learns from factual information in general. Perhaps, when we are told the answer to a question of fact, we often have an exaggerated feeling of having known it all along. In hindsight terms, we may believe that the facts we hear more or less had to be the answers to their respective questions, just as events reported to have happened seem as though they had to happen. Such a tendency could have serious implications. If we underestimate how much we are learning from the facts presented in a particular context, we should feel less reason to go on learning. If what we learn does not surprise us, then we overestimate how much we know already. Such exaggeration would be another expression of what Dawes (1976) has called "cognitive conceit."

The following studies looked for a *knew-it-all-along* effect using general knowledge questions taken from almanacs and encyclopedias. Each question had two alternative answers, one of which was correct (e.g., absinthe is [a] a precious stone or [b] a liqueur). Subjects assigned a probability of being correct to either the first or to the second answer of each question. In Experiment 1, one group of subjects (*memory*) first answered a set of such questions, then were told the correct

This research was supported by the Advanced Research Projects Agency of the Department of Defense and was monitored by the Office of Naval Research under Contract N00014-76-C-0074 (ARPA Order No. 3052) under subcontract from Decisions and Designs, Inc. to Oregon Research Institute.

Helpful comments by Barbara Combs, Bernard Goitein, Michael Hays, Sarah Lichtenstein, and Paul Slovic and the programming skills of Bernie Corrigan and Mark Layman are gratefully acknowledged, as are the thoughtful observations of two anonymous reviewers.

Requests for reprints should be sent to Baruch Fischhoff, Decision Research, 1201 Oak Street, Eugene, Oregon 97401.

answers, and finally attempted to remember their own responses. The *reliability* group first answered and then attempted to remember their responses; however, they were not told the answers. *Hypothetical* subjects saw the questions for the first time with the answers indicated. They were then asked to respond as they would have had they not been told what the answers were.

Subjects in the hypothetical group afflicted by a knew-it-all-along bias should overestimate how well they would have done on the questions had they not been told the answers. Specifically, they should believe that they would have assigned higher probabilities to alternative answers that were reported to be correct and lower probabilities to incorrect alternatives than are assigned by people who have, in fact, not been told the answers. Such overestimation would parallel hindsight subjects' tendencies to exaggerate the probability they would have assigned to reported events had they not been told what happened (Fischhoff, 1975a). A strong enough effect would interfere with memory subjects' abilities to remember their own responses, leading them to remember having been more knowledgeable than they actually were. This effect would parallel the memory distortions found by Fischhoff and Beyth (1975) with subjects asked to recall predictions that they had made 2 weeks to 6 months previously. These subjects remembered having assigned higher probabilities to events they believed to have occurred and lower probabilities to ones they believe had not occurred than was actually the case.

A knew-it-all-along effect would generalize not only the hindsight results but also those found in a number of recent studies of people's inability to ignore information they have observed or been told. Such information has included inadmissible evidence in court proceedings (Sue, Smith, & Caldwell, 1973), acts of aggression followed by evidence of mitigating circumstances (Zillman & Cantor, 1976), and behavior elicited by the manipulations of experimental psychologists that subjects are told to disregard in debriefing sessions (Ross, Lepper, & Hubbard, 1975). Detailed information about how these biases work in their most general form should improve our understanding of how knowledge is stored, altered, and retrieved. Success in ameliorating these biases, the goal of Experiment 2, could have implications for a variety of applied settings.

The questions chosen ranged in difficulty from very easy to very hard to very deceptive, reflecting items to which most, half of all, or almost no subjects knew the answers. This selection was designed to reveal, for example, whether the answers to difficult and deceptive items lead subjects to most strongly overestimate what they knew all along—or whether they produce a feeling of "there was no way I could have known that." Having subjects respond at times to the correct alternative and at times to the incorrect alternative allowed investigation of the differential effect, if any, of being told that an answer is correct and of being told that it is incorrect.

Experiment 1

Method

Design. Subjects were assigned to one of three groups: memory, reliability, and hypothetical. In Part 1 of Experiment 1, each group was asked to answer 75 questions such as "absinthe is (a) a precious stone or (b) a liqueur" by assigning a probability of being correct (between .00 and 1.00) to one of the two alternatives. Half of each group responded to the first alternative of each question and half to the second alternative. A set of unrelated tasks lasting about 1 hour separated Part 1 from Part 2.

In Part 2, all subjects received the same set of 25 test questions. For memory and reliability subjects, these 25 were a subset of the 75 questions used in Part 1. They were told, "A number of items which you just answered are repeated in the next questionnaire. Please answer each item exactly as you did before. That is, remember (or reconstruct if you have forgotten) your original responses as accurately as you can." In Part 2, memory subjects found the correct answer to each item circled "for your [the subject's] general information"; reliability subjects were not told the answers. The hypothetical group did not see the test questions in Part 1; instead, they saw 25 other questions of similar difficulty. As with the memory group, the correct answer was circled. The hypothetical group was asked to respond as they would have had they not been told

what the answer was. They were told that their responses would "enable us [the experimenters] to evaluate the perceived difficulty of these items."

Stimuli. Questions covered a wide variety of content areas such as history, music, geography, nature, and literature. Alternative answers were created to produce items of varying difficulty to elicit a full range of probability responses.

Subjects. Ninety-three paid volunteers who had responded to an advertisement in the University of Oregon student paper participated. Assignment to groups was determined by subjects' preferences for experiment time and date.

Procedure. All Part 1 questionnaires and answer sheets were collected and checked for completeness before distribution of the unrelated tasks. Completion time was approximately 40 minutes for Part 1 and 20 minutes for Part 2.

Results

Reliability. Reliability subjects' memories were quite good. They accurately remembered 66% of all Part 1 responses and 85% of their .00, .50, or 1.00 responses. There were no apparent differences in Part 1 and Part 2 responses other than slight regression effects, which were similar for both correct and incorrect responses. Thirty of the 50 alternative answers (two for each of the 25 individual questions) had identical median probabilities in Parts 1 and 2. The regression lines relating mean Part 1 (x) and Part 2 (y) responses were virtually identical for correct ($\tilde{y} = .99x + .03$; $r = .97$) and incorrect ($\tilde{y} = .89x + .06$; $r = .95$) alternatives.

Memory. Memory subjects' memories were also quite good. However, their overall proportion of correctly remembered responses (.53) was significantly ($z = 5.67$) lower than that for the reliability group (.66), indicating that being told the correct answers to the questions did interfere with their memories. Memory subjects accurately remembered less than a quarter of the original responses that were not .00, .50, or 1.00 (while correctly remembering 76% of those three response types).

In 72% of the cases in which they did not remember their original response, memory subjects recalled having assigned a higher likelihood to the correct answer than they actually had. Whether measured by the proportion of misremembered responses that constituted "increases" or by the extent of the increase, this tendency was greatest when the correct answer was unlikely. Comparing subjects' original and remembered responses revealed a mean increase of .002 when the original response was greater than or equal to .50 and a mean increase of .18 when the original response was less than .50. Apparently, subjects who were told the correct answer had difficulty remembering how they could ever have found it completely unreasonable.

Telling subjects that particular alternatives were incorrect produced similar but appreciably weaker effects. For those wrong answers originally assigned a high probability of being correct ($\geq .50$), subjects remembered having assigned lower probabilities in 76% of all cases in which the original and remembered responses were not equal. For other wrong answers, there was a strong tendency to remember having assigned higher probabilities than actually had been assigned (79% of misremembered responses). Indeed, although a significant majority of all misremembered responses ($z = 2.04$) were lower than the originals, these memory results might be interpreted as regression effects. Because no such regression was found with the reliability group, provision of the correct answers might be seen as simply reducing the reliability of subjects' responses—with incorrect items.

Hypothetical. Figure 1 compares the mean probability that hypothetical subjects believed they would have assigned to the 50 answers had they not been told which were correct, with the mean probability assigned by reliability and memory subjects in Part 1 (who actually had not been told). Hypothetical subjects substantially overestimated how much they would have known without being told the answer. As with the memory group, the effect was greater for correct than for incorrect alternatives.

For 20 of 25 items (sign test; $p = .002$), hypothetical subjects believed that they would have assigned higher probabilities to correct alternatives than uninformed subjects actually did; whereas for 15 of 25

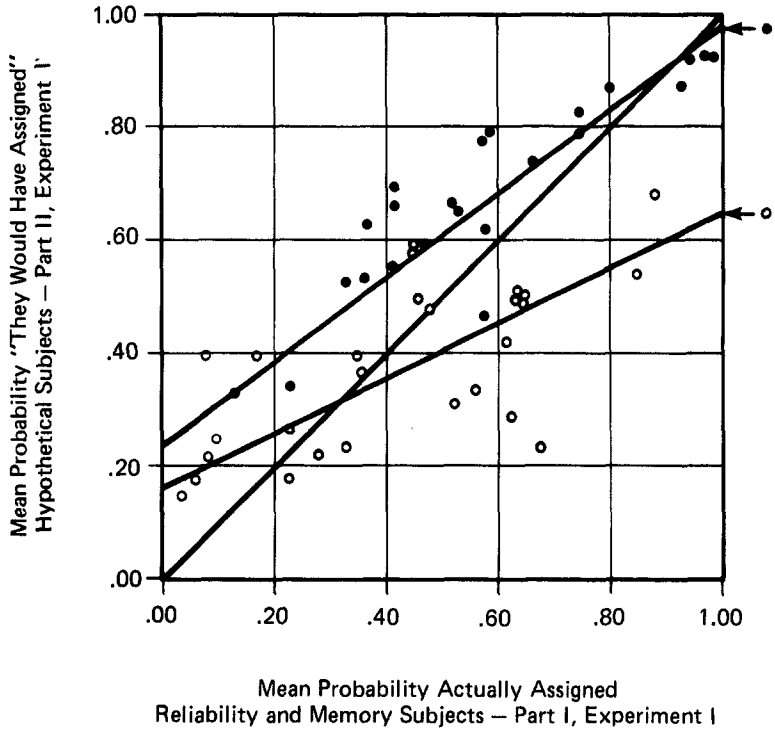

Figure 1. Item-by-item responses of hypothetical subjects in Part 2 (responding as if they had not been told the answer) and of reliability and memory subjects in Part 1 (who actually had not been told the answer) for Experiment 1. (Open circles represent incorrect alternatives; filled circles represent correct alternatives.)

incorrect alternatives (sign test; $p = .212$), they believed that they would have assigned lower probabilities. For correct alternatives, mean hypothetical probabilities were typically .10 to .25 higher for all but the most likely alternatives (where such increases were impossible). The corresponding pattern with incorrect alternatives was much less consistent. Over the 25 items, the mean size of the knew-it-all-along effect was .10 for correct alternatives (.65 vs. .55) and .05 for incorrect ones (.40 vs. .45). Table 1 summarizes these results for the three conditions in Experiment 1.

To provide some indication of individual differences, each Part 2 response was scored as *biased* if it was higher than the mean Part 1 response for that item for a correct alternative or lower than the mean Part 1 response for an incorrect alternative. The number of biased responses per subject ranged from 9 to 23 (maximum = 25) with $M = 16.3$ and $s = 3.7$.

Discussion

Apparently, people do overestimate both how much they knew (memory) and would have known (hypothetical) without being told the answers to general-knowledge questions. The hour between Part 1 and Part 2 produced numerous memory errors with reliability subjects but no systematic biases. The same memory impairment coupled with being told the answers led to a substantial bias with memory subjects. The size of the hypothetical effect, a mean change of about .10 over all questions, is similar to that found with hypothetical subjects in Fischhoff (1975a); the memory effect there was smaller than that in Fischhoff and Beyth (1975), where 2 weeks to 6 months separated prediction and memory tasks. Replication of Experiment 1 with an alternative response format (first choose the correct answer, then indicate the probability that your choice

Table 1
Mean Probabilities Assigned

Group	Correct alternative	Incorrect alternative	No. responses
Experiment 1 (25 items)			
Part 1			
Reliability and memory	.549	.447	1,700
Part 2			
Reliability	.570	.450	750
Memory	.605	.451	950
Hypothetical	.645	.396	625
Experiment 2 (27 items)			
Part 1			
Reliability and memory[a]	.526	.479	1,836
Part 2			
Hypothetical	.622	.396	1,053
Hypothetical warning	.618	.407	918
Hypothetical debiasing	.631	.403	972

[a] These are the mean responses to the 27 items of reliability and memory subjects in Part 1 of Experiment 1.

is correct) produced similar results (Fischhoff, Note 1) for the hypothetical group but a smaller memory effect. The latter may have been due to having Part 2 follow Part 1 immediately rather than after 1 hour of interpolated tasks.

Use of items varying widely in difficulty revealed that the less likely a reported answer, the greater the effect. This, too, replicated an effect that appeared to emerge from the less systematically selected items of Fischhoff (1975a) and Fischhoff and Beyth (1975). Now that it has been obtained, this result seems reasonable both because unlikely answers are more surprising and salient and, thus, should have greater impact and because the constraints on the effect imposed by the natural limits of the probability measure (.00 and 1.00) are more distant. For example, a subject told that there are actually seven and not three dwarfs in the story of Snow White (one of the undeceptive questions) has been told very little. Nor is such a subject able to believe that he or she would have assigned that answer a probability of being correct much greater than the mean assignment of .99 by subjects who were not told the answer.

At the other extreme, subjects greatly underestimated how surprised they should be at the answers to deceptive questions. Consider, for example, the following deceptive item:

Aladdin's nationality was

a. Persian
b. Chinese

Subjects in Part 1 who were not told that he was Chinese assigned mean probabilities of .13 to the correct answer (b) and .84 to the wrong answer (a). In Part 2, memory subjects remembered having assigned probabilities, the means of which were .25 and .79 for b and a, respectively; hypothetical subjects believed that they would have assigned probabilities with means of .32 and .54 to b and a, respectively.

A surprising result of Experiment 1 was the weakness of the knew-it-all-along effect with incorrect answers. Although in Part 2, memory and hypothetical subjects generally assigned lower probabilities to incorrect alternatives as hypothesized, the effect was restricted to alternatives originally assigned a high probability of being correct. For the easiest incorrect alternatives (those

assigned low probabilities in Part 1), subjects underestimated how much they knew and would have known without being told the answer. This differential effect is particularly intriguing because it resembles a similar interaction found by Fischhoff and Beyth (1975). Their subjects showed a strong tendency to remember having given higher probabilities than they actually had to reported events but a much weaker tendency to remember having given lower probabilities to events that had not happened. Before speculating on the source of the present interaction, it seems sensible to verify its existence. Therefore, Experiment 2 uses as its first group a direct replication of the hypothetical group of Experiment 1.

Experiment 2

If the knew-it-all-along effect is seen as a judgmental bias, an important applied question arises, namely, What will it take to enable people to appreciate how much they have learned from reported answers? Experiment 2 attempted to answer this question by telling a second group, the *hypothetical-debiasing* subjects, about the bias and various things they might do to reduce it. In prospect, it seemed possible that debiasing information could either achieve the desired effect or make the subjects' tasks more difficult. A third group, the *hypothetical-warning* subjects, was exhorted to work as hard as possible. This group was included as a control for the possibility that the hypothetical-debiasing group might be affected by the tone rather than the content of the debiasing instructions.

Method

Design. Three versions of the hypothetical condition of Experiment 1 were used. All subjects first responded to 75 items, assigning a probability from .00 to 1.00 to either the first or the second alternative of each. As before, the correct answer was not indicated for these Part 1 items. In Part 2, the correct answer to each of 27 additional items was circled, and subjects were asked to respond as they would have had they not been told the answer. The hypothetical group of Experiment 2 replicated the similar group in Experiment 1. The hypothetical-debiasing group was informed about the bias noted in Experiment 1 and encouraged to avoid it. As a control for the effect of heightened attention, a hypothetical-warning group (instructed to work as hard as possible) was included.

Instructions. In Part 2, all three groups were told

> On the following pages you will find a number of additional items which we intend to use in a subsequent study, identical to the one in which you just participated. Although the correct answers to these items are indicated by a circle, we would like you to respond to them as you believe you would have responded had you not been told the answer. Your responses will enable us to evaluate the perceived difficulty of these items.

For the hypothetical-warning group, the following was added:

> Your responses are extremely important to us. The effort you invest in them will largely determine the value of our subsequent study. Please devote as much attention to this task as you can. Thank you.

The hypothetical-debiasing group was also told

> On previous occasions in which we have given people this task, we have found that they exaggerate how much they have known without being told the answer. You might call this an I-knew-it-all-along effect.
>
> Consider, for example, the following question: Adaptive radiation refers to (a) evolutionary changes in animal life toward increased specialization or (b) the movement of animals to a more suitable environment for survival. A group of people who were told that the correct answer was a believed that they would have assigned a probability of about .60 to a. A group of people who were not told the answer believed that the item was a toss-up. They assigned a probability of .50 to a. Another group of people who were told that the correct answer was a believed that they would have assigned a probability of .40 to b, the incorrect answer. Again, people who were not told the answer assigned a probability of .50 to b. As you can see, people who were told the answer to an item assigned a higher probability to the correct answer or a lower probability to the incorrect answer than they might have if they had not been told the answer.
>
> In completing the present questionnaire, please do everything you can to avoid this bias. One reason why it happens is that people who are told the correct answer find it hard to imagine how they ever could have believed in the incorrect one. In answering, make certain that you haven't forgotten any reasons that you might have thought of in favor of the wrong answer—had you not been

told that it was wrong. In addition to figuring out how the correct answer fits in with whatever else you know about each topic, devote some attention to trying to see how the incorrect answer might also have fit in.

At the other extreme, however, be careful not to overcorrect and sell yourself short by underestimating how much you would have known without the answer.

Subjects. One hundred and nine paid volunteers who responded to an advertisement in the University of Oregon student paper were assigned to the three conditions according to their preference for experiment time and date. The hypothetical-debiasing group was run last to eliminate the remote possibility that word might get out about the bias.

Stimuli. Two additional items of known difficulty (taken from Part 1 of Experiment 1) were added to the 25 used in Part 2 of Experiment 1. These added items filled in gaps in the distribution of item difficulty. The order of the original 25 items was varied slightly because the random order used in Experiment 1 resulted in a disproportionate number of very difficult items toward the end of the test.

Results

Hypothetical. The responses of hypothetical subjects in Experiments 1 and 2 were generally indistinguishable. For the 25 common items, mean probability assignments from the two experiments were, respectively, .66 and .64 for correct answers and .39 and .40 for incorrect answers. The scatter plot for the Experiment 2 hypothetical group (not shown) was remarkably similar to the Experiment 1 hypothetical group shown in Figure 1. Table 2 shows that the corresponding regression lines were virtually identical. Hypothetical subjects again believed that they would have assigned higher probabilities to correct answers and lower probabilities to incorrect answers than they actually would have (see also Table 1). Again, the effect seems greatest for the most surprising answers. And again, the effect was greater and more consistent for correct than for incorrect answers.

There was a dramatic difference in the reliability of the means for the two sorts of answers. For the 25 correct alternatives used in both experiments, the means from Experiment 1 correlated .93 with the corresponding means from Experiment 2. The accompanying regression line was indistinguishable from the identity line ($\tilde{y} = .998x - .008$, where x is the Experiment 1 mean and y is the Experiment 2 mean for each answer). For the 25 incorrect answers, the correlation was merely .63 ($\tilde{y} = .54x + .19$). If these two correlations are assumed to be from independent samples, the difference is highly significant ($z = 3.02$).

Hypothetical-warning group. The hypothetical-warning group was included to evaluate the effect of exhorting subjects to

Table 2
Regression Statistics for Comparing Hypothetical Subjects with Subjects Actually Not Told the Answers

Group	Correct alternative			Incorrect alternative		
	Slope	Intercept	r	Slope	Intercept	r
Experiment 1						
Hypothetical[a]	.73	.24	.92	.49	.16	.65
Experiment 2						
Hypothetical	.74	.24	.93	.37	.22	.74
Hypothetical warning	.92	.14	.87	.54	.16	.76
Hypothetical debiasing	.75	.74	.88	.52	.16	.75

Note. Parameters are for $\tilde{y} = ax + b$, where y is the mean probability of being correct assigned to a particular answer by hypothetical subjects in Part 2, and x is the corresponding mean assignment by reliability and memory subjects in Part 1 when they did not know the answer. In Experiment 1, $df = 23$; in Experiment 2, $df = 25$.

[a] Shown in Figure 1.

work harder. As seen in Tables 1 and 2, this manipulation had little effect. Correlations between hypothetical means and hypothetical-warning means for corresponding items were .89 and .73 for correct and incorrect alternatives, respectively.

Hypothetical-debiasing group. Tables 1 and 2 present the responses of subjects told about the bias and how they might avoid it. Clearly, the debiasing manipulation failed. As Table 1 shows, the overall knew-it-all-along effect was unchanged. The scatter plot summarized in Table 2 was remarkably similar to those of the other hypothetical groups. The present means correlated .96 (correct answers) and .85 (incorrect answers) with those of hypothetical subjects (Experiments 1 and 2 combined). The only interesting results revealed by a variety of post hoc analyses were (a) a correlation of .64 between how large the original knew-it-all-along effect was and how much it was reduced or increased by the debiasing instructions—for incorrect alternatives and (b) no correlation between the amount of debiasing for the two alternatives to each question—even though with hypothetical subjects the size of the knew-it-all-along effect was highly correlated for these paired alternatives ($r = .67$).

Discussion

Experiment 2 demonstrated the robustness of the knew-it-all-along effect; it was both replicated and left unaffected by exhorting subjects to work harder or telling them to beware of bias in their responses. Certainly, there are times when one feels "I never would have known that" when told the answer to a question (e.g., "How do you say *perch* in Estonian?"). Quite possibly there were at least a few hypothetical and memory subjects for whom Aladdin's nationality came as a total surprise and was recognized as such. As a group, however, they exaggerated how much they would have known without being told about even so surprising an answer. I believe that people are capable of conjuring up a feeling of having known something about the most disparate facts. Problems arise not from being able to make sense out of just about anything (which is probably quite adaptive) but from failing to realize how much one's perceptions have been restructured by being told the answer.

These results closely parallel those in the studies of hindsight (Fischhoff, 1975a, 1975b; Fischhoff & Beyth, 1975). Of the possible explanations offered there, one seems best generalized to account for the present effects. On hearing the answer to a question (be it "What happened next?" or "Where was Aladdin born?"), people may immediately integrate that answer with whatever else they know about the topic. The purpose of this integration is to create a coherent whole out of all relevant knowledge. It may involve both reinterpreting previously held information to make sense out of it in light of the reported answer and strengthening associative links with reasons supporting the reported answer. These processes are so natural and immediate that people don't appreciate the effect that hearing the answer has had on their perceptions. As a result, they overestimate how obvious the answer appeared (memory) or would have appeared (hypothetical) before its correctness was affirmed. Even when told to do so, it is evidently extremely difficult to deprocess so important a bit of information as the right answer, inadmissible evidence, or an act of aggression followed by mitigating circumstances (see also Collins & Loftus, 1975; Kvale, 1974, 1975; Loftus, 1975).[1]

The unreliability of responses to incorrect alternatives and the differential

[1] These knew-it-all-along results might be explained as a social desirability effect, with subjects deliberately altering their Part 2 probability judgments to exaggerate the extent of their own knowledge. I find this possibility unlikely because of the nonevaluative nature of the Part 2 instructions. Memory subjects' tasks in Part 2 were described as a test of memory not knowledge. Hypothetical subjects were told that their responses would be used to judge the perceived difficulty of items to be used in other experiments. Hypothetical-debiasing subjects were directly challenged to avoid such a bias.

knew-it-all-along effect with correct and incorrect alternatives merit some consideration. The processing demands of responding to each of the two kinds of alternatives provides one possible explanation. With correct alternatives, subjects are asked, "Had you not been told that this was the correct answer, how likely would you have thought that it was?" With incorrect alternatives, subjects are asked, "Had you not known that this was *not* the correct answer, how likely would you have thought that it was?" The latter question may simply be more difficult to handle and so produces less reliable results. Jones (1966a, 1966b) has found that subjects have difficulty following instructions with implicit negatives, a difficulty they do not attempt to overcome by independently recoding the instructions to remove the negative element (see also Clark, 1969).

An alternative explanation is that being told that an answer is right has greater impact than being told that it is wrong. Subjects not told the answer to an item presumably evaluate the relative strength of the evidence supporting *both* of the possible alternatives. Subjects told the answer, however, may first figure out why the correct alternative is correct and, only then, devote some attention to why the other alternative is incorrect. Hypothetical and memory subjects using this procedure would rework their cognitive representations of correct alternatives more than those of incorrect ones. The less the reworking, the less is the effect. An analogous possibility was raised by Fischhoff and Beyth (1975) to account for the differential effect they found with events that had and had not occurred. Nonoccurrences are, in a way, nonevents leading to little restructuring of one's perceptions.

Finally, one possible explanation offered by Fischhoff (1975a) seems to merit rejection. In its most general form, that explanation contended that people estimate the likelihood that they originally assigned or would have assigned to an answer reported to be correct in two steps. First, they assign it a probability of 1.00 (its current probability of being correct); second, they look for reasons to adjust that initial value downward. For incorrect answers, the adjustment proceeds upward from .00. Inadequate adjustment, as demonstrated by Tversky and Kahneman (1974), would produce a knew-it-all-along effect. This explanation fails, however, to account for (a) why the debiasing instructions did not produce increased adjustment and reduced bias, and (b) why it should be easier to adjust upward from .00 than downward from 1.00.

More work is needed on the general problem of how people answer hypothetical questions of the form "What would I know (and think) if I had not been told ____ [some fact]?" Such research would expand our understanding of hindsight, inadmissible and mitigating evidence, debriefing, and the perceived informativeness of factual information. It would also improve our understanding of the way in which cognitive representations are restructured to include new information and of the reversibility of that restructuring. The questions to be asked are fairly obvious. Unfortunately, we may not fully appreciate how much we have learned when we manage to answer them.

Reference Note

1. Fischhoff, B. *What did you learn in school today? Not much.* Unpublished manuscript, 1976. (Available from Baruch Fischhoff, Decision Research, 1201 Oak Street, Eugene, Oregon 97401.)

References

Clark, H. H. The influence of language in solving three-term series problems. *Journal of Experimental Psychology, 1969, 82,* 205–215.

Collins, A. M., & Loftus, E. F. A spreading-activation theory of semantic processing. *Psychological Review,* 1975, *82,* 407–428.

Dawes, R. M. Shallow psychology. In J. S. Carroll & J. W. Payne (Eds.), *Cognition and social behavior.* Hillsdale, N.J.: Erlbaum, 1976.

Fischhoff, B. Hindsight ≠ foresight. *Journal of Experimental Psychology: Human Perception and Performance,* 1975, *1,* 288–299. (a)

Fischhoff, B. Hindsight: Thinking backward? *Psychology Today,* April 1975, pp. 71–76. (b)

Fischhoff, B., & Beyth, R. "I knew it would happen" —Remembered probabilities of once-future things. *Organizational Behavior and Human Performance,* 1975, *13,* 1–16.

Jones, S. Decoding a deceptive instruction. *British Journal of Psychology*, 1966, *57*, 405–411. (a)

Jones, S. The effect of a negative qualifier in an instruction. *Journal of Verbal Learning and Verbal Behavior*, 1966, *5*, 497–501. (b)

Kvale, S. The temporality of memory. *Journal of Phenomenological Psychology*, 1974, *5*, 7–31.

Kvale, S. Memory and dialectics: Some reflections on Ebbinghaus and Mao Tse Tung. *Human Development*, 1975, *18*, 205–222.

Loftus, E. F. Leading questions and the eye-witness report. *Cognitive Psychology*, 1975, *7*, 560–572.

Ross, L., Lepper, M. R., & Hubbard, M. Perseverance in self-perception and social perception: Biased attributional processes in the debriefing paradigm. *Journal of Personality and Social Psychology*, 1975, *32*, 880–892.

Sue, S., Smith, R. E., & Caldwell, C. Effects of inadmissible evidence on the decisions of simulated jurors: A moral dilemma. *Journal of Applied Social Psychology*, 1973, *3*, 345–353.

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

Zillman, D., & Cantor, J. R. Effect of timing of information about mitigating circumstances on emotional responses and retaliatory behavior. *Journal of Experimental Social Psychology*, 1976, *12*, 38–55.

Received April 21, 1976 ■