

"I Could Never Have Expected It to Happen": The Reversal of the Hindsight Bias

DAVID MAZURSKY AND CHEZY OFIR

Hebrew University of Jerusalem

[Following the stock market crash]:
". . . in my wildest dreams, I would not
have imagined this," Richard Barris,
The New York Times, October 20, 1987

It has been hypothesized and demonstrated in previous research that individuals' recall of predictive judgments is typically distorted by knowledge of the outcomes of the events predicted. This is attributed to the tendency to downgrade the surprise element associated with the outcomes and the adoption of an "I knew it all along" attitude. The present study identifies the limits of this hypothesis by showing that following the exposure to unexpected events, individuals may react by expressing an "I did not expect this to happen" response and recall predictions opposite to their judgment of the event after its occurrence. In other words, the recall of past judgments may be biased in a direction contrary to rather than consistent with subsequent judgments. Three experiments were conducted in different contexts to test the boundary conditions of the "I knew it all along" hypothesis. The findings in all three experiments suggest that following unexpected and surprising events, recall judgments are biased in a direction opposite to that predicted by the hindsight bias. © 1990 Academic Press, Inc.

After learning of the occurrence of an event, people tend to exaggerate the extent to which they had foreseen the likelihood of its occurrence. People also underestimate the effect of hindsight knowledge on their judgment. This phenomenon is known as hindsight bias (Fischhoff, 1975).

The effect of hindsight bias has been demonstrated in the context of judgment of historical, political, and social events (Fischhoff, 1975; Fischhoff & Beyth, 1975; Leary, 1981; Pennington, 1981), factual general knowledge statements (Slovic & Fischhoff, 1977), and medical diagnosis (Arkes, Wortmann, Saville, & Harkness, 1981). These findings have im-

Co-authorship is equal and the names are ordered alphabetically. The authors acknowledge the support of the Shey Foundation and the Davidson Center. They also acknowledge the Editor and two anonymous reviewers for their helpful comments on a previous draft of this paper, and Michael Ross and Yaacov Schul for their comments in the early stages of this research. Address correspondence to: School of Business Administration, Hebrew University, Mount Scopus, Jerusalem 91905, Israel.

portant implications. First, if the effect of the hindsight bias is unnoticed, it can distort people's ability to judge and study the past. Second, if in hindsight people feel that they "knew it all along," they will underestimate the informativeness of facts, and overestimate their knowledge.

The explanation proposed for this bias is that new information or outcome knowledge is immediately assimilated with what is known about the event. "The purpose of this integration is to create a coherent whole out of all relevant knowledge" (Fischhoff, 1977, p. 356). This account for the bias agrees with Loftus and Loftus (1980) who suggest that memory for complex events will be erased and updated by new information, given that it is efficient and that the new and old information are inconsistent. In providing retrospective judgments without easy access to memory for foresight knowledge, people's responses may reflect the belief that they "knew it all along" (Fischhoff, 1975). In this situation, people may employ the "anchoring and adjustment" heuristic (Tversky & Kahneman, 1974), whereby the hindsight judge uses outcome knowledge as an anchor without adjusting it downwards to a sufficient degree (cf. Fischhoff, 1977, p. 357). Updated information also enhances the "availability" of scenarios leading to the occurrence of the known outcome, potentially biasing judgments upwards (Fischhoff, 1975).

The bias can be reduced by asking hindsight subjects to explain the reported outcome along with other possible outcomes that did not occur. This procedure makes other outcomes more "available" and ". . . helps people retrieve some of the perspective which was available in foresight" (Slovic & Fischhoff, 1977, p. 55). This debiasing procedure was also demonstrated and supported in the research conducted by Davies (1987). Hasher, Attig, & Alba, (1981) indicated that under the circumstances in which the bias is generally observed, it is extremely difficult to recall prior knowledge. However, when subjects are forced to use some other retrieval plan, they may succeed in obtaining access to prior knowledge. This is achieved by discrediting outcome information. Unfortunately, as pointed out by Davies (1987), these conditions are not practical as debiasing procedures as they occur infrequently in everyday life.

Overall, the cumulative evidence consistently supports the existence and robustness of hindsight bias. The only known exceptions reported in the judgment literature are either when outcome information is discredited or when hindsight judges are asked to consider and explain other outcomes which did not occur.

A more natural boundary of this bias can be hypothesized on the basis of studies on social attribution. This line of research suggests that unexpected outcomes may lead people to generate reasons and explanations to account for the inconsistency between the old and the new information. Lau and Russell (1980), for example, when examining newspaper articles

covering major sporting events, found that the number of explanation attempts elicited by unexpected events was larger than the number of explanations elicited by expected events (cf. Lau, 1984). In other studies, in which subjects were provided with a description of a person performing actions that are congruent and others that are incongruent with his/her character, incongruent behaviors were reported to be best recalled (Crocker, Binns, & Weber, 1983; Srull, 1981). Using a similar experimental paradigm, Clary and Tesser (1983) exposed subjects to a general description of an actor and then provided additional information consistent or inconsistent with the original description. When asked to retell the story they had read, those subjects who were initially exposed to an expectancy-disconfirming story more often provided excuses, justifications, and unread explanations in the retelling task than those receiving an expectancy-confirming story (see also Pyszczynsky & Greenberg, 1981).

Recently it has been suggested that relatively high recall in this context results from "special processing" (Hastie & Park, 1986). This processing involves attempts to explain the unexpected incongruent behaviors (Hastie, 1984). Based on a review of these and other studies, Weiner (1985) concluded that unexpected events elicit a causal search to explain the outcome. After experiencing an unexpected or surprising event, people attempt to search for causal explanations. Since it is often hard to recall expectations, especially for relatively complex events, people may reconstruct their memories. Under these conditions they may feel "we did not expect this to happen." Thus, when asked to recall their expectations, people tend to exaggerate the difference (rather than the consistency) between foresight and hindsight expectations.

The present study aims at investigating the effect of unexpected events on people's recalled expectations. In the experimental conditions (i.e., hindsight conditions) of Experiment 1, the outcome information was manipulated to give different levels of "unexpectedness." These groups were compared with a control (i.e., foresight) group. It was hypothesized that in the face of a surprising outcome, expectations recalled in hindsight will shift in a direction that contrasts with outcome information.

EXPERIMENT 1

Experiment 1 involved participants' assessments of one of two movies. Both movies exemplified marketing methods and skills, the difference between them being in the level of quality. The movies were shown to first-year business students, a relevant audience for this task. We hypothesized that subjects assessing the high-quality movie (judged by three raters and the previous year's students as representing very high quality) will recall having assigned lower levels of expectations for such movies than subjects exposed to the low-quality movie.

Method

Subjects

Eighty-four students enrolled in undergraduate business courses participated in the study. They were told that “we want to learn your opinion about movies as teaching aids in studying business-related topics.” The study was conducted during a class period. Approximately equal numbers of males and females participated in the study.

Stimulus Material

Three independent raters selected two movies out of a stock of about 30 available movies. The first movie selected was judged as representing high quality and the second as representing low quality in terms of their potential for exemplifying real life business applications and their usefulness as a complementary educational method. In addition, the high-quality movie was recommended to us by a colleague who had used it with a previous year’s student audience.

Procedure and Design

Subjects were randomly assigned to three conditions, one control and two experimental. Subjects in the control group were handed a questionnaire aimed at assessing their expectations from movies as teaching aids. This group did not watch either of the movies prior to filling out the questionnaire. Experimental subjects were randomly assigned into two separate rooms where the movies were shown. The high-quality movie was screened only for its first 30 min (and stopped at a stage that was not felt as an unnatural end by the subjects) to match the 30-min length of the low-quality movie.

Following the screening, all subjects in the two experimental conditions filled out a two-part questionnaire, one part relating to their recalled expectations (compatible with the control group’s questionnaire), and the other relating to the movie itself and subjects’ reactions to it.

Dependent Measures

Control group questionnaire. Nine questions pertained to expectations from movie use as a teaching aid. These included beliefs and assessments of such movies along different dimensions (e.g., exemplification of current business methods and representing improvement in learning about business realities). All questions were rated on 7-point scales (with anchors (1) very unlikely to (7) very likely).

Experimental groups’ questionnaire. The first part of this questionnaire was compatible with the control group’s questionnaire except that it required subjects to recall their expectations from movies as a teaching aid

prior to watching the film. Finally, an overall evaluation of the movie ((1) very poor to (7) very good) was measured to confirm the raters' evaluation of the movies as representing high and low quality.

Results

Preliminary analysis indicated that the mean of the performance measure was significantly higher for the high-quality movie (mean = 5.32 and mean = 3.78 for the high- and low-quality movies, respectively, $t(54) = 3.84, p < .001$).

In order to generate the dependent variable for the main analysis, an index of preexposure expectations (i.e., the judgments made by subjects in the control group) was obtained by combining the nine expectation ratings with equal weights ($\alpha = .72$, Cronbach, 1951). This procedure was repeated for the experimental groups ($\alpha = .82$).

The levels of preexposure expectations with their recalled postexposure counterparts were in line with the hypothesis (i.e., lower recalled expectations in the high-quality group) (for the high quality mean = 5.31, for the low quality mean = 5.74, and for the control group mean = 5.77). A one-way ANOVA was performed on the three levels of expectations (i.e., preexposure expectations and recalled expectations for the low- and high-quality movies). The analysis revealed significant between-group variation ($F(2,81) = 3.48, p < .03$).

A comparison between the control and low-quality movie condition revealed no significant difference ($F(1,81) < 1$). However, a comparison between the control and the high-quality condition indicated that recalled expectations of the group assessing the high-quality movie were significantly lower than the expectations of the control group ($F(1,81) = 5.58, p < .02$). This analysis suggests that the between-group variation stems primarily from the lower value of post-exposure recall in the high-quality movie condition.

Discussion

In line with much of the theoretical and empirical evidence suggesting that outcome information is likely to induce memory reconstruction, Experiment 1 indicated that subjects' levels of recalled expectations were affected by the exposure to the two movies each representing different levels of quality. Unlike many previous studies, however, the direction of this reconstruction was one of overestimation rather than underestimation of the difference between the evaluation of the present experience and recalled expectations. The mediocre responses to the low-quality movie did not differ significantly from recalled expectations (in fact, recalled expectations of this group were almost identical to the preexposure expectations in the control group). On the other hand, judgments made by

subjects in the high-quality movie condition come with lower levels of recalled expectations.

Experiment 2 was conducted to replicate the findings of Experiment 1 and extend them by incorporating measures of surprise and disconfirmation of expectations. These measures were posited to assess postexposure realization that a change had indeed occurred. Accordingly, we hypothesized that a high degree of surprise will be associated with reconstruction of expectations to reflect the feeling of "I did not expect it to happen."

EXPERIMENT 2

Method

Subjects

Thirty-six students participated individually in the study. They were randomly assigned to three treatment groups. Subjects were paid for their participation. An approximately equal number of males and females participated in the study.

Stimulus Material

To assure a relatively low predictive baseline and "surprise" at the outcome, we chose a product generally perceived as poor in performing the task it is claimed to perform. Based on a pilot study conducted among 180 students, we selected a suction hook as the appropriate product. (A suction hook is a device composed of a metal hook attached to a plastic cone which is pressed onto glass, ceramic, or any smooth surface. It is commonly used for hanging towels or light-weight clothes.) The judgments obtained in a separate group of subjects who rated the hook only on the basis of past experience and knowledge confirmed the selection of the product as a low-credibility device.

The second criterion motivating the selection of this product related to the experimenter's control over product performance. By perforating a tiny hole in the plastic cone (which could not be noticed by the subject), the hook would fall after being attached for less than 1 min. Such a hook was considered defective. All other hooks held a weight of 15 lb for more than 1 min and, thus, were considered "good." In selecting the six hooks to be used in the six successive trials (see below), the experimenter had full control over the desired combination of good and defective hooks. This manipulation masked the true intention of the study.

Procedure and Design

The study was presented as a marketing project being conducted by a company which had recently introduced a new brand of a suction hook.

This hook was claimed to be manufactured from a new plastic material which distinguished it from other available hooks.

The experiment consisted of three between-subjects condition. In all three conditions the experiment involved hanging a 15-lb weight on six suction hooks, each for 1 min. For each treatment, a different combination of good and defective hooks was prepared before the subject entered the laboratory. Subjects were, of course, blind to the combination of good and defective hooks. In the poor-performance condition, subjects were handed a basket containing six defective hooks; in the intermediate condition, equal numbers of good and defective hooks were selected; and in the superior performance condition, all six hooks were good.

Dependent Measures

Upon completing the six successive trials, subjects in all groups were handed a questionnaire composed of two sets of questions. The first set pertained to subjects' past, inquiring about recalled expectations and past satisfaction with this type of product. The second set pertained to post-trial assessment of product performance, satisfaction, and surprise. In the first set, two dependent measures were included, one relating to past behavior (whether subjects had used such hooks in the past)¹ and the other relating to satisfaction with suction hooks in the past with the range: (1) not satisfied at all to (7) extremely satisfied.

The second set consisted of three questions: (1) satisfaction with the present brand (using the same scale as for recalled satisfaction); (2) the disconfirmation of product performance expectations (on a 7-point scale with five anchor points ranging from "Worse than expected" to "Better than expected"); and (3) surprise inherent in product performance (with anchors compatible with the disconfirmation question).

Results

Posttrial Surprise and Disconfirmation

Preliminary examination of the data revealed that subjects in the high- and intermediate-performance conditions were more surprised (mean = 4.42, 5.92, 5.92, respectively) and expressed higher disconfirmation (mean = 4.42, 6.17, 6.17, respectively) than those in the poor-performance condition. The pattern of results appears to be in line with the manipulation except that no differences exist between the intermedi-

¹ Since satisfaction with suction hooks in the past is a main variable in the experiment, subjects had to have some previous experience with such hooks. Experience with similar hooks was therefore used as a screening variable. In line with this procedure, the first 12 qualifying subjects in each condition were included in the experiment.

ate and the high conditions, possibly due to a ceiling effect. The omnibus F -tests, however, are significant in both analyses suggesting that both surprise level ($F(2,33) = 11.1, p < .001$) and disconfirmation ($F(2,33) = 12.5, p < .001$) were significantly lower in the low-performance condition.

Recalled vs Posttrial Judgments

A mixed 3×2 ANOVA was performed with one between-subjects factor (three levels of product performance) and one within-subjects factor (recalled vs posttrial satisfaction; see Fig. 1). Posttrial satisfaction was higher than recalled satisfaction ($F(1,33) = 20.71, p < .001$). The main effect due to level of performance was not significant.

More importantly, the interaction effect was significant ($F(2,33) = 4.36, p < .02$) suggesting that better performance of the hooks was associated with an increased gap between posttrial and recalled satisfaction. Analysis of the simple effects supports this interpretation; no significant difference was obtained in the "defective only" condition ($F(1,11) < 1$). However, the difference in the "mixed" condition ($F(1,11) = 4.81, p <$




FIG. 1. Satisfaction as a function of product performance and judgment reference.

.06) and in the "good only" condition ($F(1,11) = 11.2, p < .01$) was more pronounced.

Performance evaluation and disconfirmation of expectation increased with the increase in the proportion of high-quality hooks in the experimental task. This increase, however, was not obtained in the case of recalled expectations. On the contrary, subjects tended to derogate past evaluations when exposed to unexpected performance outcome. In Experiment 2 as in Experiment 1, the response to unexpected performance was manifested both by assigning higher ratings to the performance measure and by derogating past expectations. Experiment 3 provides a replication of these findings in a different domain.

EXPERIMENT 3

Method

Subjects

Fifty students were seated in front of a personal computer. They were recruited and tested individually and were compensated for their participation.

Procedure and Design

Subjects were randomly assigned to two experimental conditions. They were presented with a description of a graphic software and were asked to perform a trial task on the personal computer. After completing the task, subjects were asked to indicate their pretrial quality expectations of the graphic software.

The trial task involved producing a particular type of graph using a program which was new to all the participants. The software was capable of producing high-quality graphs by utilizing a richer set of options than those offered by other available programs.

The set of options offered to the subjects during the trial task was experimentally manipulated between subjects: one experimental group used the program without any limitations; the other experimental group was exposed to versions of the program with a limited number of options which produced graphs of poorer quality. Upon completing the task, the two groups were handed a three-part questionnaire, one section pertaining to subjects' pretrial quality evaluations of the graphic software (measuring eight beliefs about the software quality), the second section pertaining to posttrial quality ratings of the software. These ratings were measured along scales compatible to the pretrial measures. Finally, subjects were asked to rate the degree to which they were surprised by the program's quality and to rate the degree to which the experience either

confirmed or disconfirmed their expectations regarding the quality of the software.

Results

The main dependent variable for the analysis is based on the pretrial and posttrial quality evaluations of the graphic software. The measure was obtained by summing-up the eight beliefs regarding the software quality. An alpha coefficient computed for this index indicated acceptable reliability (pretrial $\alpha = .78$, and for the posttrial $\alpha = .90$).

The mean quality judgments of the various conditions are presented in Fig. 2. As hypothesized, a significant interaction was obtained ($F(1,48) = 8.07, p < .01$) from a mixed two-factor ANOVA with level of quality as the between-subjects factor (high vs low quality) and judgment reference as the within-subjects factor (hindsight: recall of pretrial evaluation vs posttrial evaluation). A significant main effect for the judgment reference factor was obtained as well ($F(1,48) = 5.41, p < .03$). Simple effects




FIG. 2. Quality evaluations as a function of the graphic software quality and judgment reference.

computed for each of the two quality levels indicated no significant difference between recalled expectations and posttrial evaluations for the low-quality group ($F(1,24) < 1$) but a significant difference for the high-quality condition ($F(1,24) = 14.24, p < .001$).

As expected, the subjects in the high-quality condition were significantly more surprised (mean = 5.88) than subjects in the low-quality condition (mean = 4.79) ($t(46) = 2.8, p < .01$). Similarly, subjects in the high-quality condition indicated (mean = 4.52) significantly higher disconfirmation than those in the low-quality condition (mean = 4.37; $t(46) = 2.22, p < .03$).

GENERAL DISCUSSION

It has been repeatedly demonstrated that after the outcome of an event becomes known, there is a distortion in the recall of expectations. People tend to feel that "they knew it all along," and they underestimate the effect of outcome information on their judgments. The literature in the field supports the existence and robustness of this bias which is considered to extend to historical judgment in general (Fischhoff, 1980).

The present study identifies the limits of this bias by showing that following exposure to a surprising or unexpected event, individuals may react by expressing an "I did not expect this to happen" response. Since it is hard to recall expectations, people tend to reconstruct their memory so that their judgments are biased in a direction opposed to presently held knowledge. People attribute the surprise to their inability to have foreseen an outcome such as the one obtained, and recall expectations that are lower than those actually expected. This finding is consistent with the notion that justifications and explanations are elicited following an unexpected outcome. This attempt to explain an unexpected event leads to an exaggerated adjustment in a direction opposite to the hindsight bias. In other words, present knowledge serves as an anchor and to accommodate the surprise people "recall" significantly lower expectations.

Acknowledgement of surprise in the present study distinguishes it from prior research. Previous studies have generally postulated and demonstrated empirically that outcome knowledge is not so surprising in hindsight (e.g., Slovic & Fischhoff, 1977). Fischhoff, in one of his studies (1977, p. 356), did raise the possibility of recognized surprise and indicated that this may have happened to few subjects in his experiments. As a group, however, he indicated that they behaved in a manner supporting the hindsight bias. In contrast, in the present study (Experiments 2 and 3) surprise and disconfirmation judgments were high after the outcome became known. Consequently, hindsight judgments of past expectations were negatively related to the surprising quality and performance of the outcome.

It is interesting to note a line of research in social psychology which resembles the findings regarding hindsight bias. When people are asked to recall past attitudes and behavior after unnoticed attitude change, they tend to be committed to their newly formed attitudes by selectively recalling arguments which are in favor of these new attitudes (Ross, McFarland, Conway, & Zanna, 1983; Lydon *et al.*, 1983) or by eliciting reasons and explanations that account for the current attitude.

In the above area of research, suggestions somewhat similar to those presented here are advanced. According to this perspective, the reconstruction of memory relating to past behavior and attitudes is largely influenced by characterization of the past as being either different from or the same as the present. Insofar as the theories of stability and change which are invoked maintain that no change has occurred (even if the judgment is biased), a consistency in past and present is most likely to predominate. If, on the other hand, the present is perceived as being different from the past, reconstruction of the past may overestimate change. This latter tendency was demonstrated empirically in the setting of a self-improvement program (Ross & Conway, 1986). Accordingly, poorer expectations may be "recalled" if people desire a change while in actuality there has been little or no change between past and present status. The present study has shown that the "exaggerated change" bias may be obtained under more natural circumstances in which no motivational factors are posited to account for the findings.

Finally, it could be argued that a halo effect of the measurement instrument contributed to the results of studies such as those reported here. For example, in the case of exaggerated consistency, the bias may be attributed to the instrument employed to assess both the performance and the recalled judgments, rather than representing a genuine error. In the case of exaggerated change, the interpretation of the results based on such a rival explanation is unlikely. If at all, halo would undermine the effects obtained in the reported experiments. Likewise, the settings in which the present experiments were conducted involved minimal interface with experimenters, thus precluding the possibility of demand characteristics in the interpretation of the data.

REFERENCES

- Arkes, H. R., Wortmann, R. L., Saville, P. D., & Harkness, A. R. (1981). Hindsight bias among physicians weighing the likelihood of diagnosis. *Journal of Applied Psychology*, *66*, 252-254.
- Clary, E. G., & Tesser, A. (1983). Reactions to unexpected events: The naive scientist. *Personality and Social Psychology Bulletin*, *9*, 609-620.
- Conway, M., & Ross, M. (in press) Getting what you want by revising what you had. *Journal of Personality and Social Psychology*.

- Crocker, J., Binns, D., & Weber, R. (1983). Person memory and causal attribution. *Journal of Personality and Social Psychology*, *44*, 55–66.
- Cronbach, L. G. (1951). Coefficient alpha and the internal structure of tests. *Psychometrika*, *16*, 297–334.
- Davies, M. F. (1987). Reduction of hindsight bias by restoration of foresight perspective: Effectiveness of foresight-encoding and hindsight-retrieval strategies. *Organizational Behavior and Human Decision Processes*, *40*, 50–68.
- Fischhoff, B. (1980). For those condemned to study the past: Reflections on historical judgment. In R. A. Shweder (Ed.), *New directions for methodology of social and behavioral science*, *4*, San Francisco: Jossey-Bass.
- Fischhoff, B. (1977). Perceived informativeness of facts. *Journal of Experimental Psychology: Human Perception and Performance*, *3*, 349–358.
- Fischhoff, B. (1975). Hindsight-foresight: The effect of outcome knowledge on judgment under uncertainty. *Journal of Experimental Psychology: Human Perception and Performance*, *1*, 288–299.
- Fischhoff, B., & Beyth, R. (1975). "I knew it would happen"; Remembered probabilities of once-future things. *Organizational Behavior and Human Performance*, *13*, 1–16.
- Goethals, G. R., & Reckman, R. F. (1973). The perception of consistency in attitudes. *Journal of Experimental Social Psychology*, *9*, 491–501.
- Hasher, L., Attig, M. S., & Alba, J. W. (1981). I knew it all along: Or, did I? *Journal of Verbal Learning and Verbal Behavior*, *20*, 86–96.
- Hastie, R., & Park, B. (1986). The relationship between memory and judgment depends on whether the judgment task is memory-based or on-line. *Psychological Review*, *93*, 258–268.
- Hastie, R. (1984). Causes and effects of causal attribution. *Journal of Personality and Social Psychology*, *46*, 44–56.
- Lau, R. R. (1984). Dynamics of the attribution process. *Journal of Personality and Social Psychology*, *46*, 1017–1028.
- Lau, R. R., & Russell, D. (1980). Attributions in the sports pages: A field test of some current hypotheses in attribution research. *Journal of Personality and Social Psychology*, *39*, 29–38.
- Leary, M. R. (1981). The distorted nature of hindsight. *Journal of Social Psychology*, *115*, 25–29.
- Loftus, E. F., & Loftus, G. R. (1980). On the permanence of stored information in the human brain. *American Psychologist*, *35*, 409–420.
- Pennington, D. C. (1981). The British firemen's strike of 1977/78: An investigation of judgment in foresight and hindsight. *British Journal of Social Psychology*, *20*, 89–96.
- Pyszczynski, T. A., & Greenberg, J. (1981). Role of disconfirmed expectancies in the investigation of attributional processing. *Journal of Personality and Social Psychology*, *40*, 31–38.
- Ross, M., McFarland, C., Conway, M., & Zanna, M. P. (1983). Reciprocal relation between attitudes and behavioral recall: Committing people to newly formed attitudes. *Journal of Personality and Social Psychology*, *45*, 257–267.
- Ross, M., & Conway, M. (1986). Remembering one's own past: The reconstruction of personal histories. In R. M. Sorrentino & F. T. Higgins (Eds.), *Handbook of motivation and cognition* (pp. 122–144). New York: Guilford Press.
- Slovic, P., & Fischhoff, B. (1977). On the psychology of experimental surprises. *Journal of Experimental Psychology: Human Perception and Performance*, *3*, 544–551.
- Srull, T. K. (1981). Person memory: Some tests of associative storage and retrieval models. *Journal of Experimental Psychology: Human Learning and Memory*, *7*, 440–563.

- Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, *185*, 1124–1131.
- Weiner, B. (1985). Spontaneous causal thinking. *Psychological Bulletin*, *97*, 74–84.
- Wong, P. T. P., & Weiner, B. (1981). When people ask “why” questions, and the heuristic of attributional search. *Journal of Personality and Social Psychology*, *40*, 650–663.
- Wood, G. (1978). The “I knew it all along” effect. *Journal of Experimental Psychology: Human Perception and Performance*, *4*, 345–353.

RECEIVED: March 20, 1986