JOURNAL OF VERBAL LEARNING AND VERBAL BEHAVIOR 20, 86–96 (1981)

I Knew It All Along: Or, Did I?

LYNN HASHER, MARY S. ATTIG, AND JOSEPH W. ALBA

Temple University

Other studies have reported that when subjects are presented with outcome feedback they are unable to remember their original knowledge state (the "knew-it-all-along effect"). In the present studies, feedback was followed by manipulations which were intended to invalidate it. In the first experiment, we failed to discredit the feedback and so report a knew-it-all-along effect under circumstances different from those reported elsewhere. In the second experiment, the discrediting instructions were successful and the effect was disrupted. Contrary to previous interpretations, the latter results indicate that feedback information is not automatically assimilated and that people can access their prior knowledge state, if the circumstances require.

Recall inaccuracies have always been of considerable empirical interest to psychologists for the insight they provide into the operation of memory. One recently reported systematic memory inaccuracy is of considerable practical interest as well. This is the "knew-it-all-along effect" (Fischhoff, 1975) in which adults appear to be unable to remember their original knowledge state (e.g., Aladdin, of the magic lamp story, originated in Persia) after the experimenter has provided them with outcome information (it was really China). The knew-it-all-along effect appears to be the product of basic memory processes rather than of demand characteristics such as an implicit instruction of the experimenter to use the outcome information just provided. Indeed, attempts to eliminate this and other demand characteristics of the task as the source of subjects' failure to recollect original knowledge have not succeeded (Fischhoff, 1977; Wood, 1978).

The explanation offered for the knew-it-all-along effect is that outcome knowledge is immediately assimilated with prior knowledge (e.g., Fischhoff, 1975, p. 279) leaving no residual information in memory about the original knowledge state. This explanation is one that meshes nicely with current assumptions about the mechanisms underlying other interesting memory inaccuracies, including the effect of leading questions asked of eyewitnesses to accidents and crimes (Loftus, 1979). Loftus and Loftus (1980) propose that memories for complex events are updated by new information and that old information will be erased (by a substitution process) if two conditions obtain: (a) if it is efficient to do so; and, (b) if the earlier information is inconsistent with the updated information. These criteria can be directly applied to the circumstances that prevail in the hindsight or knew-it-all-along literature. If one learns the outcome of some historical event (e.g., who won the conflict between the Gurkas and the British) or the truth about some fact (e.g., Aladdin originated in China), it would be inefficient to maintain one's earlier incorrect or uncertain knowledge and certainly illogical to believe two mutually exclusive facts (e.g., that Aladdin originated in both China and Persia).

These update and erase notions about the functioning of memory also mesh easily with theories which predict other memory inaccuracies based on the role of schemas in the encoding of stories, conversations, and other complex social situations. These

This research was supported by Grant 1R01 MH33140-01 from the National Institute of Mental Health. We thank Henny Sanft, James Arbuckle, and Edward Gracely for their help.

0022-5371/81/010086-11$02.00/0

structures (schemas, story grammars, contexts, frames) are thought to influence the type and extent of information actually stored in memory about any given event (see Friedman, 1979, for a brief review of such theories). Insofar as memory is less than accurate, the possibility of fabrications guided by a schema (as in reconstruction or default values) always exists.

However, some recent work with prose materials suggests that memory may contain far more detailed information about complex events than schema theories would predict (e.g., Baker, 1978; Kintsch & Bates, 1977; Alba, Alexander, & Hasher, Note 1) and that while these details may be extremely difficult to retrieve they can be produced under at least some circumstances (Hasher & Griffin, 1978). The Hasher and Griffin study is of special interest here because it demonstrated that even after a week's retention interval, manipulations introduced at the time of retrieval could alter the nature *and extent* of recall of a prose passage. In that experiment, subjects read a story constructed to conform equally well to two different themes (e.g., a hunter and a convict). While reading the story, individual subjects believed it to be about one or the other of the two possible themes. Patterns of recall of the stories suggested that while some subjects engaged in what may well be the typical retrieval pattern for complex materials, that is, in reconstruction, others appeared to use the presumably more difficult retrieval pattern of reproducing the story. This latter group imported fewer theme-relevant ideas than did the former. What is critical for present purposes, however, is that the "reproducing" group recalled more of the ideas actually present in the story than did the "reconstructing" group. This difference in retrieval patterns was introduced by a manipulation which, it was argued, led some subjects to disregard the most easily accessible retrieval cue in that situation, the major theme or topic of the passage. Forced to use another retrieval plan, these subjects were able to recall information apparently inaccessible to subjects who did use the major theme as their retrieval cue.

With these findings in mind, we attempted to alter the retrieval plans of subjects in a knew-it-all-along procedure. Our goal was to demonstrate that subjects in such a situation can indeed remember their original knowledge state. Such a finding would be important because it would establish limits to the knew-it-all-along effect as well as to the extent to which assimilation processes are believed to operate in memory.

Experiment 1

The general procedure we used to produce the knew-it-all-along effect was based on that of Wood (1978). All subjects first rated on a 7-point, true-to-false scale, a set of plausible statements pertaining to a variety of knowledge domains (e.g., the arts, current affairs, sports). A second rating trial followed after a brief retention interval. Between the two rating trials, experimental subjects were provided with feedback about half the items they had rated. They were instructed to ignore this information on the rerating trial. In the critical experimental condition, the feedback was followed by a manipulation (loosely based on that of Hasher & Griffin, 1978) intended to invalidate or disconfirm the feedback. If subjects can disregard the feedback, their second ratings should look like their initial ratings, as well as like the ratings of a group of control subjects who received no feedback. Such a finding would support the contention that original knowledge state information is retained in memory. If feedback is, as Fischhoff argues (1975, p. 279), automatically assimilated with prior knowledge, the two experimental conditions should show second rating scores that are more in keeping with the feedback than with their original knowledge, a demonstration of the knew-it-all-along effect.

It should be noted here that the idea that current information is automatically assimilated to previous knowledge and leaves no trace of this knowledge is an implicit

theory of forgetting which has considerable popularity (e.g., Gentner & Loftus, 1979; Glass, Holyoak, & Santa, 1979, p. 71). The general notion is that new information which is relevant to previously stored information will activate that information and stand in its stead. This could be called an update-and-erase theory of memory.

Method

Procedure. This was a three-phase experiment. In the first and third phases, subjects rated each of a series of statements on a 7-point scale ranging from 1 indicating positively false to 7 indicating positively true. The statements rated on the two occasions were identical except for the order in which they occurred. Ratings were self-paced.

In the second phase, two groups of subjects were given feedback about information for half of the statements they had just rated. This was done by allowing subjects 3 minutes to look over a list of statements divided into two groups of 10. The experimenter indicated which set was true and which was false. The experimenter suggested that the subject might "want to study the statements for a few minutes."

The critical group in the experiment then received a brief manipulation intended to preclude the likelihood that the outcome information would be used on the subsequent rerating task. The experimenter took back the typed feedback sheet, casually looked it over, seemed taken aback, and then said to the subject, "Oh no, did I tell you that the first ten (statements) were true? I should have said the last ten." This implied of course that the statements originally identified as true were really false and those identified as false were really true. The experimenter also said that this was not a serious mistake for purposes of the experiment because the idea was for them to rerate the statements as they had done initially. Subjects did not get to look over the list of statements after that instruction.

A third group of subjects was tested in a no-feedback control situation. In the 3 minutes in which experimental subjects were given information about some statements, subjects in the No-Feedback Control condition were asked to solve the nine-dot problem. This is a rather difficult task (e.g., Weisberg & Alba, in press) in which subjects have to connect, without lifting their pens, three rows of three dots arranged in a square, using just four straight lines.

On the second rating trial, all subjects were asked to rerate the items as they had the first time. Experimental subjects were told to ignore the feedback information and try to remember their original ratings. Control subjects were asked to remember their original ratings.

Materials. A total of 40 statements referring to plausible but obscure facts in a variety of knowledge domains were selected from an original set of 140 used by Hasher, Goldstein, and Toppino (1977) as well as by Wood (1978). Twenty true and twenty false statements were selected from norms generated by college students who were asked to rate statements on the same 7-point scale used in the present experiment. The items selected to serve in the present study had mean ratings that ranged from 3.5 to 4.5. This criterion was chosen to allow for the opportunity to see rating changes induced by the feedback.

The 40 statements were divided into two sets of 20 and, during the feedback interval, experimental subjects were given feedback about one of the two sets of statements. Within each of these two sets of 20 there were three unique arrangements of statements which on feedback were arbitrarily called "true" or "false." With three subsets of "true" vs "false" statements for each of two unique sets of statements on which feedback could be given, there were six different arrangements of materials. Within each of the three conditions, equal numbers of subjects were assigned to each of the six sets of materials.

Subjects. The subjects were university students tested in small groups. There were 36 per condition.

Results and Discussion

Each subject in the experimental conditions contributed six means to the analysis, one for each of the three item types (called true, called false, and no feedback provided) on each of two rating trials. Each subject in the No-Feedback Control condition contributed two means, one for items on trial 1 and one for items on trial 2. In addition each subject contributed scores which represented the difference between his or her mean rating on trial 2 and on trial 1 for each item type. The means of these values may be seen in Table 1.

An initial set of analyses compared all item types in all conditions on the initial rating trial. There were no significant differences among items, F's < 1, indicating that the items serving in all conditions were equally plausible to all groups of subjects. Thus any effects observed as a result of feedback manipulations are unlikely to be attributable to item differences.

The second analysis also considered just the items on which no feedback was given. It assessed whether there was any change in ratings introduced by the events occurring in the short retention interval (approximately 4 minutes) between rating trials 1 and 2. As is apparent from Table 1, the values given on the two rating trials are nearly identical—for all three groups of subjects. This observation was confirmed by an ANOVA which revealed no significant effects for trials, conditions, or item types, or for any of their higher-order interactions, all F's < 1.

Thus subjects who receive no feedback at all show no change in their ratings across the two trials. Subjects who are given feedback on some items (even if that feedback is subsequently dismissed as incorrect by the experimenter) rerate items on which they received no feedback much as they did the first time. The values the experimental subjects assign are extremely similar to those assigned by the control subjects. Thus the feedback manipulations appear not to influence judgments on statements

TABLE 1
MEAN RATINGS ON TRIALS 1 AND 2 AND DIFFERENCES BETWEEN TRIALS 2 AND 1

	Standard feedback			Disconfirmed feedback			No feedback		
Item type	Trial 1	Trial 2	T2 – T1	Trial 1	Trial 2	T2 – T1	Trial 1	Trial 2	T2 – T1
"True"	3.90	4.29	.39	4.06	3.40	−.66			
"False"	4.04	3.55	−.49	3.94	4.17	.23			
Control	4.10	4.09	−.01	4.13	4.20	.08	4.08	4.10	.02

for which no information is provided. Memory for original ratings appears to be excellent.

The next set of analyses concerns just the two experimental conditions. The dependent measure for these analyses was, for each item type, the difference between the mean rating on trial 2 and that on trial 1. A positive value indicates that ratings increased (and, given that they started at the "uncertainty" point of the scale, moved toward the true end of the scale) between the two trials; a negative value indicates that ratings decreased (and so moved toward the false end of the continuum). An ANOVA comparing the three item types for both conditions revealed a significant conditions × items interaction, $F(2, 140) = 22.04$, $MS_e = .66$. Because of this interaction, the important effects are more easily seen if each condition is discussed separately.

Consider first the performance of subjects in the standard feedback condition. Items called true by the experimenters show an increase in ratings from trial 1 to trial 2, items called false show a decrease in ratings, and the ratings for the no-feedback items show no change. These pairwise differences are all significant (Newman–Keuls test). This is of course the pattern of results termed the "knew-it-all-along effect" and the findings are a straightforward replication of Wood's experiment (1978). Subjects' ratings closely follow the information provided by the experimenter even when subjects are instructed to ignore the feedback and to remember their original ratings. This pattern of results supports the update and erase view of the impact of new information which is relevant to something already stored in memory.

The performance of the subjects in the disconfirmed feedback condition also reveals a difference among items. Here again, the items on which no information is provided show no change between the two sessions. The pattern of change in ratings for the other two item types is now quite different from that seen in the standard feedback condition. Here, items initially called true show a decline across the test trials and so move toward false. Items initially called false show an increase and so move toward true. A Newman–Keuls test (at .05) showed that the "true" items declined significantly relative to the no-feedback items. While the "false" items showed an increase in ratings between the two trials, this increase was not significantly greater than the "change" seen for the no-feedback items. The subjects in the disconfirmed feedback condition are clearly not doing what we had originally anticipated; they do not disregard the experimenter's feedback and remember their original ratings. Rather, it appears as if these subjects update their knowledge with the initial feedback and then update it again when informed of the inadvertent reversal of the true and false labels.

Rather than eliminating the knew-it-all-along effect, we appear to have demonstrated that the tendency to update memory by incorporating new, relevant information is extremely strong; subjects are not disrupted by a sequence of contradictory information. They appear to treat the disruption as an honest error and to treat the feedback (once corrected for the experimenter's error) as valid. The tendency to update memory is also highly item specific; subjects appear to be able to keep three sets of statements (true, false, and no feedback) segregated. And, finally, the tendency to update memory operates extremely quickly; subjects in the disconfirmed feedback condition received the reversal information immediately after reading the statements and at that time did not have the opportunity to restudy the statements. Nonetheless, their judgments moved in the direction of the "corrected" feedback. Thus we appear to have evidence that strengthens rather than weakens the notion that a memory program rapidly updates or assimilates new information leaving the

person with no access to his or her prior knowledge state.

If updated information does supplant previous information, it does so with high levels of accuracy based upon the evidence, suggesting the keen ability of subjects in the disconfirmed feedback condition to segregate statements into categories depending upon whether or not feedback was provided, and if so, its nature. It was this evidence for accuracy, together with a reanalysis of the likely impact of our disconfirming manipulation which persuaded us not to abandon our original hypothesis. We had intended to persuade subjects to disregard our feedback. To this end we told them that a mistake had been made, the true and false categories were interchanged. It is possible (in fact, likely, based on the outcome of the second experiment) that the systematic reversal of the true and false categories in a plausible context ("I accidentally reversed them") enabled subjects to continue to believe in the credibility of the feedback and so to continue to use this easily available information during the rerating trial. Thus we made a second attempt to create circumstances under which feedback would be disregarded.

EXPERIMENT 2

> *They* [eyewitnesses to a crime] *need to be shocked into total recall. They're not seeing anything that happened with fresh eyes.* (P. D. James)[1]

Suppose it is the case that after a memory has been updated a person discovers that either the source of that updating lacks any credibility or that the information provided is actually incorrect. If no past record is preserved in memory, such a person would be in an extremely difficult position if forced to supply information. She or he would have two choices: (a) to admit ignorance; or, (b) to supply known misinformation. A third alternative is possible—if original knowledge is preserved in memory—exert the effort and use alternative retrieval strategies to gain access to that knowledge state. In this experiment, two feedback conditions were designed to be analogous to this scenario.

Methods

The three phases of Experiment 1 were repeated with standard experimental subjects receiving feedback between the two rating sessions and no-feedback control subjects solving a problem during that interval. Two new experimental conditions were tested. The critical manipulations were again introduced just after the original feedback was taken back from the subject. In one case (called Disconfirmed: Wrong), the experimenter baldly said that "the answers that I just gave you on the feedback sheet were wrong. Now I want to see if you can ignore those answers and rate the statements exactly as you did before." The experimenter offered no justification for the error. In the other experimental condition (called Disconfirmed: Mistake) the experimenter said, "It looks to me as if these feedback sheets have a lot of mistakes on them. Some of the ones listed here as true I'm sure are false; and some listed here as false I'm sure are true. I'm not sure how it happened but someone must have made a mistake. I'm sorry but it's really okay for the rest of the experiment because what I want to see is if you can ignore those answers and rate these statements exactly as you did before."

With the exception of instructional manipulations just mentioned, all procedures and materials in this experiment were identical to those used in Experiment 1. Because of limitations on the numbers of subjects available, only 18 subjects were tested in the two replication conditions (standard feedback and no-feedback control), while 36 subjects were run in each of the two experimental conditions.

Results and Discussion

The mean ratings for each item type may be seen in Table 2. The same analysis plan

[1] James, P. D. *Shroud for a nightingale*. New York: Popular Library Edition, 1971. P. 99.

TABLE 2
MEAN AND STANDARD DEVIATIONS[a] OF RATINGS ON TRIALS 1 AND 2 AND DIFFERENCES BETWEEN TRIALS 2 AND 1

Item type	Standard feedback			Disconfirmed: Wrong			Disconfirmed: Mistake			No feedback		
	Trial 1	Trial 2	T2 – T1	Trial 1	Trial 2	T2 – T1	Trial 1	Trial 2	T2 – T1	Trial 1	Trial 2	T2 – T1
"True"	3.98	4.36	.37	3.91	3.94	.03	4.04	4.29	.24			
	(.64)	(.84)		(.60)	(.74)		(.61)	(.84)				
"False"	4.07	3.62	−.46	4.08	3.91	−.17	3.97	3.87	−.10			
	(.45)	(.72)		(.56)	(.61)		(.59)	(.63)				
Control	3.97	4.05	.08	3.92	3.94	.02	4.08	4.06	−.01	3.99	4.01	.02
	(.58)	(.50)		(.49)	(.45)		(.34)	(.41)		(.40)	(.45)	

[a] Standard deviations in parentheses.

used in the first experiment was followed here. An initial check on differences among the various items on the first trial revealed no significant differences, F's < 1. Again, items which will serve in the various conditions appear to be rated about the same by all subjects before the experimental manipulations are introduced.

A second analysis was conducted on just those items for which no feedback was provided. This analysis compared trial 1 with trial 2 ratings for all four conditions. No significant effects were found, F's < 1, indicating that there was no change across the two rating trials when the subjects received no feedback. Again, this is the case whether, as in the case of the control condition, subjects received no outcome information at all, or whether, as in the case of the experimental conditions, subjects received outcome information—but on other items.

The next set of analyses was conducted on the change scores shown by subjects in the three experimental conditions. A positive value means that subjects' belief in the statements increased across the trials; a negative value, that belief decreased. There were significant effects for items, $F(2, 174) = 17.60$, as well as an interaction between items and instructional conditions, $F(4, 174) = 2.996$, $MS_e = .243$. The effects are best understood by considering each instructional group separately.

The performance of subjects in the standard instructional condition once again demonstrates the knew-it-all-along effect: The values assigned to the no-feedback items show almost no change. Belief in items the experimenters called true increased while belief in statements called false decreased. A Newman–Keuls test, $p = .01$, shows that each of these pairwise differences is significant.

The issue of greatest interest here is whether either of the two types of disconfirming instructions eliminated the knew-it-all-along effect. Consider first the condition in which subjects were told after reading the feedback that there were mistakes in

the listings. A simple effects test on the three item types revealed significant differences among them, $F(2, 70) = 5.45$, $MS_e = .21$. It is the case that relative to the no-feedback control items those called true show an increase in their ratings and those called false show a decrease. However, a Newman–Keuls test shows that, at the .01 level, only the extreme difference between the true and false statements was significant; at the .05 level, the difference between items called true and controls becomes significant. The difference between the control and false statements is not significant even at the .05 level. This is certainly not the robust pattern of the effect as seen in the standard conditions of Experiments 1 and 2.

Consider next the condition in which subjects were told after reading the feedback that that feedback was wrong. The change scores here, for all items, are smaller than we have seen them before and a simple effects comparison revealed no significant effects among them, $F(2,70) = 2.06$, $MS_e = .23$, $p = .13$. Subjects in this condition do not show the knew-it-all-along effect: Their ratings for the informed items do not differ from their ratings for items about which they receive no information. Since these latter do not show any change in ratings from trial 1 to trial 2, the inference is that neither do the ratings for informed-about items. A conservative test of this is available from a repeated measures analysis of trials 1 and 2 performance on the three item types for just these subjects. This analysis reveals no significant effects.

Thus it appears to be possible to reduce the knew-it-all-along effect (at least for the false items in the Mistake condition) and also to entirely eliminate it (as in the Wrong condition). Subjects can, in highly unusual circumstances, rerate items as they originally did even when they've learned new information about those statements. These findings do not fit easily into a view of memory which suggests that one automatic control process is the update of old information coupled with its complete erasure.

We argue here that subjects can, under some circumstances, gain access to their original knowledge—and that it is memory for their original knowledge state which accounts for their ability to rerate items as they did initially. Two alternative interpretations, both based on the assumption that subjects do not remember their original knowledge, can be considered. One alternative suggests that subjects adopt a decision rule which says always choose 4, the value that represents "I don't know." Since the statements were originally selected by the experimenters from norms which suggested that the mean ratings would be 4 on the initial trial (which was indeed the case), the use of such a decision rule on the rerating trial would make it appear as if subjects were remembering their original knowledge when in fact they were not. This alternative interpretation would suggest that the variability of ratings on the postfeedback trial should be smaller than on the prefeedback trial, when the subject is not using the "I don't know" rule. A glance at the standard deviations for ratings (Table 2) on trials 1 and 2 for the true and false items is especially revealing: In no instance was the variability on the rerating trial smaller than on the initial trial. Whether feedback was disconfirmed (as in the experimental conditions) or not (as in the standard feedback condition), the variability of judgments increased.

A second interpretation of the values on the rerating trials suggests that subjects apply a random response decision rule, first on trial 1 and then again on trial 2, thus producing the same mean rating values for the sets of items. An assessment of the degree of relationship between ratings on the two trials can be obtained from correlations calculated for individual subjects for items on trial 1 and trial 2. The mean correlations for subjects in the standard, mistake, wrong, and no-feedback conditions are .71,

.78, .72, and .84, respectively. The degree of relationship is far too high to make an independence argument viable.[3]

Thus the fact that reratings for two conditions look like the original ratings cannot be attributed to a systematic subject response bias—either to apply the same value for each item in a set or to respond randomly. Rather, it appears as if subjects can in fact remember their original knowledge, when the proper retrieval environment exists.

General Discussion

A major legacy of Bartlett's (1932) work in contemporary cognitive psychology is the concept of assimilation. This is the notion that new information will be incorporated into existing knowledge structures in such a way as to lose unique identifying tags. A unique tag might be a marker which would enable the subject to identify some piece of information as inferred rather than given, or as imagined rather than seen, or as suggested rather than real. Frequently cited examples of such assimilative processes abound: subjects will under some circumstances integrate into a cohesive structure a series of overlapping statements so as to be little able to distinguish the original from a synthesized version (Bransford & Franks, 1971); concurrently activated general knowledge plays a major role in the encoding of new information into memory (Bransford & Johnson, 1972); subjects confuse their general knowledge with newly presented information (Sulin & Dooling, 1974); eyewitnesses to accidents and crimes can be easily misled to either supplement or transform their recollection of a critical event (Loftus, 1979).

The practical implications of a memory which functions so as to assimilate related pieces of knowledge without leaving identifying markers are profound. Decisions of juries, testimony of witnesses, evaluations of the importance of new information or even the "newness" or surprise value of new information, or of the true state of one's own knowledge would be seriously undermined by the operation of this basic cognitive process.

It is therefore of considerable importance that the limits of assimilation be understood. There are two important questions here: (1) Is assimilation automatic? (2) What about the prior knowledge state remains in memory? With respect to the notion that assimilation (or updating) is automatic—in the sense of being inevitable and virtually simultaneous with the provision of new information—Wood (1978) has already supplied us with a preliminary answer: No, it is not. While Wood consistently found the knew-it-all-along effect across materials, subjects, and instructions, the size of the effect was not uniform. In particular, instructions given after feedback at the time of the second or critical rating ("remember what you said before" versus "judge these as you think your peers would") influenced the size of the effect.

The more important question though for both practical and theoretical reasons is: What is left behind in memory about the original knowledge state? Is the answer, "nothing?" Clearly not. The second experiment reveals that under some circumstances, subjects can indeed remember their prior knowledge. If subjects believed that experimenter-provided feedback was wrong they did not use their updated memories (or, rather, memory which included the updated information). Rather, they were able to gain access to their previous knowledge state and so to give approximately the same answers they had

[2] A third alternative interpretation of the similarity between values on trials 1 and 2 has also been suggested. This idea is based on the assumption that updating has occurred and prior knowledge is indeed lost. The second ratings are like the first because subjects attempt to answer the problem "afresh." No data directly apply to this suggestion. However, it seems to us to be inelegant to propose that subjects remember their ratings for items given no feedback and rework their answers to items given feedback.

given before the feedback was provided. Updating—or the provision of relevant information interpolated between one event and its recall—*does not erase the original knowledge.*

To consider what updating does do we point to two mechanisms which might explain the present findings. It seems likely (based on the present results as well as those of Hasher & Griffin, 1978) that subjects must exert unusual effort to retrieve pre-updated information. Effort is expended because subjects know that following their original retrieval plan would be absurd: in the present experiment subjects knew they had been given unreliable information; in the Hasher and Griffin study subjects knew at retrieval that a story they had read was not actually about the originally advertised topic. "Shock" treatment appears to induce subjects into expending special effort to remember an original event accurately.

It should be noted that simply asking subjects to try hard to remember will not necessarily improve accuracy (e.g., Fischhoff, 1977). What is likely to be required, in addition, is that retrieval follow a route other than it ordinarily would follow. The effect of discrediting what would be an easily retrievable piece of information (here, the most recent information about an item; in Hasher & Griffin, the title of the story) might be analogous to the impact of forcing subjects to consider negative evidence—which appears to be less easily available than positive evidence (Koriat, Lichtenstein, & Fischhoff, 1980).

The view taken here is that while initial memories may be difficult to access, they are nonetheless available—to use the original Tulving and Pearlstone (1966) terms. This distinction was then extended to the context of a retroactive interference experiment (Tulving & Psotka, 1971) which, after all, the knew-it-all-along procedure mimics (as does the Loftus leading question technique); subjects are given two successive events (the original one and the outcome one) and then their memories for the original are tested. In the classic list-learning situation, the temporary dominance of the most recent information is a well-established phenomenon. The first list is not forgotten; it is difficult to access.

Embedded in the present data is the fact that in the knew-it-all-along procedure, it is far easier to disrupt a subject's belief about statements the experimenters call false than those they call true. In fact, there was no statistically significant knew-it-all-along effect for items called false in any of the experimental conditions (Disconfirmed, Mistake, or Wrong). In both Fischhoff's work (1977) and Wood's (1978), the size of the knew-it-all-along effect is smaller for information called false during feedback than for information called true. The existence of this bias might itself suggest that old information is not erased when new information is provided.

The possible bias against believing information called false is interesting in itself for scientists and decision makers. Indeed it may well be related to the fact that evidence which disconfirms a viewpoint one adheres to (e.g., a scientific theory) tends not to be given very strong weight (Lachman, Lachman, & Butterfield, 1979).

REFERENCES

BARTLETT, F. C. *Remembering: A study in experimental and social psychology.* Cambridge, Eng.: Cambridge Univ. Press, 1932.

BAKER, L. Processing temporal relationships in simple stories: Effects of input sequences. *Journal of Verbal Learning and Verbal Behavior,* 1978, 17, 559–572.

BRANSFORD, J. D., & FRANKS, J. J. The abstraction of linguistic ideas. *Cognitive Psychology,* 1971, 2, 331–350.

BRANSFORD, J. D., & JOHNSON, M. K. Contextual prerequisites for understanding: Some investigations of comprehension and recall. *Journal of Verbal Learning and Verbal Behavior,* 1972, 11, 717–721.

FISCHHOFF, B. Hindsight ≠ foresight: The effect of outcome knowledge on judgment under uncertainty. *Journal of Experimental Psychology: Human Perception and Performance,* 1975, 1, 288–299.

FISCHHOFF, B. Perceived informativeness of facts.

Journal of Experimental Psychology: Human Perception and Performance, 1977, 3, 349–358.

FRIEDMAN, A. Framing pictures: The role of knowledge in automatized encoding and memory for gist. *Journal of Experimental Psychology: General,* 1979, 108, 316–355.

GENTNER, D., & LOFTUS, E. F. Integration of verbal and visual information as evidenced by distortions in picture memory. *American Journal of Psychology,* 1979, 92, 366–375.

GLASS, A. L., HOLYOAK, K. J., & SANTA, J. L. *Cognition.* Reading, Mass.: Addison–Wesley, 1979.

HASHER, L., & GRIFFIN, M. Reconstructive and reproductive processes in memory. *Journal of Experimental Psychology: Human Learning and Memory,* 1978, 4, 318–330.

HASHER, L., GOLDSTEIN, D., & TOPPINO, T. Frequency and the conference of referential validity. *Journal of Verbal Learning and Verbal Behavior,* 1977, 16, 107–112.

KINTSCH, W., & BATES, E. Recognition memory for statements from a classroom lecture. *Journal of Experimental Psychology: Human Learning and Memory,* 1977, 3, 150–168.

KORIAT, A., LICHTENSTEIN, S., & FISCHHOFF, B. Reasons for confidence. *Journal of Experimental Psychology: Human Learning and Memory,* 1980, 6, 107–118.

LACHMAN, R., LACHMAN, J. L., & BUTTERFIELD, E. C. *Cognitive psychology and information processing; An introduction.* Hillsdale, N.J.: Erlbaum, 1979.

LOFTUS, E. F. The malleability of human memory. *American Scientist,* 1979, 67, 312–320.

LOFTUS, E. F., & LOFTUS, G. R. On the permanence of stored information in the human brain. *American Psychologist,* 1980, 35, 409–420.

SULIN, R. A., & DOOLING, D. J. Intrusions of a thematic idea in retention of prose. *Journal of Experimental Psychology,* 1974, 103, 255–262.

TULVING, E., & PEARLSTONE, A. Availability versus accessibility of information in memory for words. *Journal of Verbal Learning and Verbal Behavior,* 1966, 5, 381–391.

TULVING, E., & PSOTKA, J. Retroactive inhibition in free recall: Inaccessibility of information available in the memory store. *Journal of Experimental Psychology,* 1971, 87, 1–8.

WEISBERG, R. W., & ALBA, J. W. An examination of the alleged role of "fixation" in the solution of several "insight" problems. *Journal of Experimental Psychology: General,* in press.

WOOD, G. The knew-it-all-along effect. *Journal of Experimental Psychology: Human Perception and Performance,* 1978, 4, 345–353.

REFERENCE NOTE

1. ALBA, J. W., ALEXANDER, S. G., & HASHER, L. *The role of semantic structures in the encoding of information.* Manuscript in preparation, 1980.

(Received April 15, 1980)