Journal of Personality and Social Psychology
1985, Vol. 49, No. 2, 338-346

Copyright 1985 by the American Psychological Association, Inc.
0022-3514/85/$00.75

The Transparency of Denial: Briefing in the Debriefing Paradigm

Daniel M. Wegner and Gary F. Coulton
Trinity University

Richard Wenzlaff
University of Texas at Austin

This research tested a new conceptualization of the impression perseverance effect. Here, as in earlier studies, some actor and observer subjects were given false feedback about the actor–subjects' performance in the experiment and then were informed during *debriefing* that the feedback had not been genuine. Other subjects, however, received a *briefing* about the falsity of the feedback before the task performance. These briefed subjects, like the debriefed subjects, subsequently made estimates of the actors' actual performance on the task that were significantly influenced in the direction of the false feedback. The briefed subjects did not, however, follow the debriefed subjects in making ability attributions to the actor in line with their performance estimates. These results cast doubt on the notion that attributional processing of the false information, as observed in the debriefing condition, is a necessary component of the perseverance effect. The idea that denied information and the denial may contribute independently to subsequent impressions is offered as an alternative explanation of briefing and debriefing phenomena.

"I deny it!" said the March Hare.
"He denies it," said the King: "Leave out that part."
(Lewis Carroll, *Alice's Adventures in Wonderland*)

Even by royal decree, it can be difficult to "leave out" information that is known to be false. Certainly, when experimental subjects are told that the information they received in a study was false, they may continue to base judgments on it nonetheless. This phenomenon was observed by Walster, Berscheid, Abrahams, and Aronson (1967), and has since been incorporated into a general model of *impression perseverance* by Ross, Lepper, and their colleagues (e.g., Ross, Lepper, & Hubbard, 1975; Ross, Lepper, Strack, & Steinmetz, 1977). According to this model, people fail to adjust their impressions sufficiently when they encounter information that discredits the evidence on which the impressions are based. The model asserts that this failure occurs when people have formed an attribution or explanation for the observed evidence. In essence, it is argued that the attribution remains as a basis for inference even when the evidence is discredited. Our purpose in this research was to weigh this model against another kind of analysis, one based on the idea that people see through denials.

Attribution and Perseverance

According to the perseverance model of Ross et al. (1975, 1977), people try to interpret or make attributions for the information they receive during an experiment. On finding that someone has failed at an experimental task, for example, both that person and observers might be likely to infer that the person has little ability. On finding that someone has succeeded, in turn, attributions of greater ability might be made. These ability attributions are inferences about the person, not characterizations of the observed events of failure or success, and they therefore might not be challenged directly by subsequent information indicating that the success or failure

We wish to thank Suzanne Anderson, Debbie Dickman, Robert Esquenazi, Nancy Newman, and Karey Sucher for help in conducting the research, and Toni Giuliano and Robin R. Vallacher for valuable comments on an earlier draft of this article.

Richard Wenzlaff is now at the University of Texas at San Antonio.

Requests for reprints should be sent to Daniel M. Wegner, Department of Psychology, Trinity University, 715 Stadium Drive, San Antonio, Texas 78284.

feedback had been false. So, although debriefing may discount the false information, subjects' interpretations and attributions remain—to form the basis for continued inferences consistent with the information.

This attributional model has been widely accepted in social research as an explanation for the instances in which experimental debriefing does not fully disabuse subjects of the impressions they formed during an experiment (Jelalian & Miller, 1984). Moreover, the model has been offered as an account of the instances outside experiments in which people seem to cling to beliefs despite the discrediting of evidence on which the beliefs were founded (Nisbett & Ross, 1980). The general adoption of this model is perplexing, however, in light of the evidence available. One kind of evidence, for example, comes from studies attempting to increase the degree to which subjects form interpretations or attributions. Instructing subjects to make attributions sometimes intensifies perseverance (Anderson, Lepper, & Ross, 1980; Ross et al., 1977), but more often has no detectable impact (Anderson, 1983; Carroll, 1978; Fleming & Arrowood, 1979; Jennings, Lepper, & Ross, 1981). Quite simply, experimental attempts to enhance perseverance by stimulating the presumed mediator of the phenomenon have not been uniformly successful.

A related strategy is to prevent subjects from making attributions by distracting them after the false feedback has been delivered. Although this tactic appears to reduce perseverance (Fleming & Arrowood, 1979), it also can be understood as a way of preventing subjects from thinking about the impression at all (see Carroll, 1978) and so cannot be taken as firm evidence for the attributional mediation of the effect. Indeed, experimental demonstrations of the impact of explanation outside the debriefing paradigm (e.g., Gregory, Cialdini, & Carpenter, 1982; Sherman, Skov, Hervitz, & Stock, 1981) are similarly indirect in their support. Because attributional processing may also make the event more available in memory (Tversky & Kahneman, 1973), most attempts to enhance or undermine explanation have been irrelevant to the question of whether explanation alone produces belief.

A different line of evidence on the role of attribution in perseverance comes from several demonstrations of variations in perseverance effects that occur independent of variations in subjects' opportunities for attributional processing. In the presence of circumstances that would seem to encourage subjects to form detailed attributions—for instance, significant delays between impressions and debriefing—perseverance effects have not been found (e.g., Hatvany & Strack, 1980; Tennen & Gillen, 1979). And perseverancelike effects have been found, in turn, even when subjects are instructed to ignore information immediately following their exposure to it (e.g., Thompson, Fong, & Rosenhan, 1981).

Taken together, these findings do not invalidate the attributional model. There is considerable evidence that attribution has effects under certain conditions—although it is unclear whether these effects are traceable to attribution per se or to the enhanced processing and availability of information that usually accompanies attributional inference. The findings observed to date, however, do seem to indicate that attribution may not be the most powerful variable at work in the debriefing paradigm. Perseverance effects seem to appear and disappear independent of attributional processing, and it is difficult to maintain at this point that perseverance is primarily traceable to attribution.

Denial Transparency

What then leads people to persevere? Our approach to this question is based on the idea that impression perseverance is inherent in the person's understanding of denied information. We believe that people can see through denials. On encountering denied information, the person typically has that information available for processing despite the denial. This is because people, unlike computers, have no "reset button" that can completely eradicate memory. Rather, people process information cumulatively, always adding to their store of knowledge, and cannot use one item of information to delete another. Even when people are directly instructed to forget information, they may restrain themselves from retrieving it in the current circumstance but still have it available for later retrieval (Geiselman, Bjork, & Fishman, 1983). After all, Freud (1925/1961) counted denial as only a weak defense mechanism for

the very reason that it seldom succeeds in erasing a thought from a person's mind. Bateson (1968), too, has argued that because a denial is a metacommunication, the denial and underlying communication are of different logical types, never truly to be combined.

Our conception of denial transparency is based on this general realization, and consists of two interrelated ideas. The first is that a denial is typically processed as an addendum to information. When people are asked to reason logically with denials, for example, attempting to solve syllogisms that involve statements such as "X is not true," they often make errors (Wason & Johnson-Laird, 1972). People are often confused by double or even single negatives, but have no comparable difficulty dealing with double positives, triple positives, and beyond. This occurs because an affirmation or a positive can be disregarded in the further processing of information; a statement such as "This is an apple" is functionally equivalent to the affirmation that "It is true that this is an apple." A denial or negation, however, represents an added step in any subsequent reasoning. The "not true" modifier must be attached to the information and carried with it at all times for rational operations on the information to be carried out. Far from erasing an impression, then, a denial *accompanies* an impression.

The second component of our conception of denial transparency is this: Because people store in memory both the initial information and its denial, questions relevant to the domain of the denied information will be answered by a cognitive process that (a) consults the initial information, and (b) attaches, with variable reliability, a denial to the answer. In essence, this formulation suggests that an impression forms a core memory that participates in subsequent processing. The denial is peripheral to this core, participating in various ways at times, but primarily as an afterthought rather than as a prelude. So, for example, when a person reads a book to find on the last page a disclaimer indicating that the contents of the book are wholly false, no thoroughgoing change reverberates through the person's cognitive representation of the book's contents. Rather, a "not true" addendum is attached to the memory trace of the book. The same sort of addendum would be attached if the disclaimer were encountered at the outset. In either case, however, when the person is then called on to answer questions about the book's topic, the person may often remember to attach the "not true" feature to the answers—and often may not. Because denied information requires that extra step, it will not be retrieved as reliably as undenied information. In particular, it will often be retrieved in its initial, undenied form.

Evidence favoring this view of denial comes from a variety of experimental settings. The notion that a denial is cognitively stripped away from an impression, for example, is substantiated in a study by Snyder and White (1981). These investigators asked subjects to test hypotheses that were formulated as denials (e.g., the hypothesis that a person is *not an extrovert*). Now as a rule, when people seek to test an affirmative hypothesis (e.g., that a person is an extrovert), they look for evidence that would confirm the hypothesis (Snyder & Swann, 1978; Wason & Johnson-Laird, 1972). However, the subjects in this study abandoned this hypothesis confirmation strategy, seeking information to indicate that the person was an extrovert even when directly instructed to test the hypothesis that the person was *not* an extrovert. Although the investigators somehow construed this finding as consistent with a confirmatory bias interpretation, their results indicate to the contrary that subjects seek to *dis*confirm denials by seeking out evidence that would confirm the impression. Here, then, people seem to process denials by looking right through them to the underlying information (see also Semin & Rosch, 1981).

Perhaps the most direct demonstration of denial transparency occurred in research conducted by Wegner, Wenzlaff, Kerker, and Beattie (1981). These studies of the impact of media innuendo indicate, for example, that when people read a newspaper headline asserting that "Bob Talbert is not linked with Mafia," they develop unfavorable impressions of Bob Talbert. Although their impressions are still not as negative as ones drawn from an incriminating assertion (e.g., "Bob Talbert is linked with Mafia"), they are nonetheless substantially less favorable than impressions derived from unincriminating assertions (e.g., "Bob Talbert arrives in City"). Such innuendo

effects in impression formation have been observed in several experiments (Wegner, 1984), each of which indicates that a simple denial is typically understood to convey instead an impression consistent with the underlying information. These effects are not significantly reduced when the denials are attributed to highly credible sources, suggesting that even an easily believed denial can lead people to accept some version of the information that it is intended to discredit.

The denial transparency model seems to indicate that denials are never very successful. We wish to point out, however, that denials can fulfill their intended logical function in certain cases. When, for instance, denials can be cognitively integrated with impressions to form *reversed* impressions, their impact may be more lasting. On hearing someone described as "not warm," for example, perceivers might cognitively transform this to "cold" under certain conditions and thereby avoid the further cognitive work of carrying along both the impression and the denial in memory (Johnson-Laird & Tridgell, 1972; Wason, 1965). Such transformations are not always logically correct, however; a person who is not warm, in fact, might also be not cold. Nevertheless, such reversals do seem to occur at times, especially when the impression has a semantic opposite available in the language. We would expect that denials that are reconceptualized in this or other ways to become affirmative statements would have a much more consistently logical impact on subsequent judgments.

Yet other exceptions to denial transparency might occur when people have little use for the information contained in the denial. When we learn that "Roy did not hold up a filling station," for instance, our lack of other information about Roy makes the hold-up salient indeed. We might well be swayed by denial transparency to think of Roy as a thief because the denied information forms our only judgmental anchor regarding Roy's nature (see Kahneman & Tversky, 1973). If we knew that Roy was extremely rich, however, or that he was a pacifist, or even that he had arrived at the filling station during the robbery, we would have much less call to draw on the denied information. Certainly, it would no longer be our only resource for judgments in the domain and so would be less available when questions about Roy were posed. In addition, many of the other facts we knew might be inconsistent with the impression (e.g., Roy could not both arrive at the station during the robbery and be the robber) and so serve as cues to remind us to retrieve the denial along with the impression. In short, although denial transparency could be a common effect in judgments based on a denial alone, it should be much less prevalent when additional domain-relevant information is available.

Denial transparency seems to be a more straightforward interpretation of the perseverance effect than the attributional processing explanation. With this in mind, we designed the present experiment to determine whether perseverance might occur under conditions that would preclude an attributional interpretation, but that would follow from the denial transparency model. Such conditions exist when subjects are *briefed* rather than *debriefed* about the false feedback they receive during an experiment. Told in advance that the information they will obtain is false, subjects would have little call to make interpretations or attributions about the false information. After all, attributional inference is not an automatic consequence of information intake (Smith & Miller, 1983), and it seems unlikely that such effortful cognitive elaboration would follow on the receipt of admittedly false information. However, briefed subjects would still be exposed to the information and its denial as in the usual debriefing paradigm. If the perseverance effect were to occur under these circumstances, it would constitute evidence favoring the denial transparency hypothesis.

Method

Overview

We designed this study to replicate in most respects the methods used by Ross, Lepper, and Hubbard (1975, Experiment 2). As in their experiment, actor and yoked observer subjects participated in a study calling for the actor subjects to discern true from false suicide notes. And, as in this prior study, the actors were randomly assigned to receive false feedback indicating either success or failure. In the present research, though, an additional variable was included in the design. Half of the participants in the aforementioned conditions were debriefed as in Ross et al.; they were told after the completion of the judgments that the feedback had been false. The remaining participants were briefed; they were told prior to the

occurrence of the judgment trials that the upcoming feedback would be false.

Subjects and Procedure

Ninety-two undergraduate students (51 females and 41 males) served as subjects in return for extra credit in their introductory psychology classes at Trinity University. They were randomly assigned to serve as actors or observers under the stipulation that no actor–observer pair should be acquaintances. The observer was ushered into the experimental room first and seated so as to be able to observe the actor unobtrusively through an opening between a pair of movable partitions. The experimenter explained that the observer's task was to watch as the other subject participated in a discrimination task, but that the other subject would not know this. The observer was instructed to record the actor's verbal responses on the task and to rate the actor's apparent degree of confidence in each response.

At this point, the actor was brought into the room and seated at a desk approximately 3 m from the observer such that the observer would have a side view of both the actor and the experimenter. The actor was told that the observer was completing a questionnaire that was part of a different experiment. The experimenter then indicated to the actor that the study was arranged to examine physiological responses during decision making. Recording electrodes were attached to the actor's fingers, and the experimenter explained that he would give the actor a series of pairs of suicide notes (taken from Schneidman & Farberow, 1957). The actor's task for each would be to discern which was a real note, one written by an actual suicide victim. The actor was instructed to read each pair, indicating his or her judgment by saying its identifying letter aloud (i.e., *A* or *B*). The experimenter explained that he would be providing immediate feedback on each trial, indicating to the actor whether the judgment had been correct or incorrect.

For subjects in the briefing condition, the experimenter explained to the actor at this point that the feedback would be false. He noted that

> I will give you feedback on your performance on each trial by saying either 'correct' or 'incorrect.' However, this feedback will not be genuine. The feedback you will receive has been predetermined and will not reflect your actual performance on this task. This procedure is used in an attempt to identify how false feedback affects physiological responses during decision making.

Actor subjects in all conditions then were administered the judgment task, receiving either success feedback (correct on 24 of 25 trials) or failure feedback (correct on 10 of 25 trials), and were asked subsequently to sit quietly and relax for five minutes "while further physiological readings are taken." For subjects in the debriefing condition, following the feedback trials the experimenter explained to the actor that

> Now that the experiment is over I can inform you that the feedback that you were given was not genuine. The feedback you received was predetermined and did not reflect your actual performance on this task. This procedure was used in an attempt to identify how false feedback affects physiological responses during decision making.

The observer was taken to a different room at this point, and both actor and observer were given questionnaires containing the key dependent measures. As in Ross et al. (1975), subjects were asked to estimate (a) how many of the 25 trials the actor actually judged correctly; (b) how many of the 25 trials the average student would judge correctly; (c) how many trials the actor would judge correctly if given a different series of 25 trials; and (d) how capable the actor is at performing the judgment task, as rated on a 7-point scale.

Results

The results for the four measures are presented in Table 1. A separate analysis of variance (ANOVA) was conducted for each measure to examine the impact of performance feedback (success versus failure) and denial timing (briefing versus debriefing), with the actor–observer variable included as a repeated measures factor.

Subjects' estimates of the actual number of successes the actor achieved were significantly influenced by the feedback manipulation, $F(1, 42) = 37.04$, $p < .001$. No other effects or interactions approached significance in this analysis. Subsequent tests of simple main effects ascertained that significant effects of feedback were present both in the debriefing condition, $F(1, 42) = 25.28$, $p < .01$, and in the briefing condition, $F(1, 42) = 14.38$, $p < .01$. Means in the success feedback condition did not differ significantly between briefing and debriefing, and means in the failure feedback condition also did not differ between briefing and debriefing. In sum, briefing and debriefing had essentially equivalent effects, leading neither actors nor observers to forsake the feedback as a cue to the actor's performance.

Estimates of the number of successes possibly achieved by an average student on this task were similarly affected only by the feedback manipulation, $F(1, 42) = 12.02$, $p < .01$. The estimate for the average student represents a baseline from which impressions of the actor would be expected to depart, were success and failure performance impressions of the actor particularly strong. With this in mind, it is worth noting that when a difference score was calculated as the estimated actual successes for the actor minus the estimated average student's successes, the significant impact of feedback was again observed, $F(1, 42) = 5.90$, $p < .05$. The actor was seen as more inclined to exceed the

Table 1
Measures of Feedback Acceptance by Experimental Condition

Measure	Briefing		Debriefing	
	Success	Failure	Success	Failure
	Actors			
Estimated initial number correct	17.50	13.33	16.50	12.10
Estimate for average student	14.75	13.25	14.50	12.70
Predicted future number correct	17.00	14.58	16.17	14.50
Rated ability at task	4.42	4.67	4.92	4.35
	Observers			
Estimated initial number correct	17.00	14.17	17.08	12.20
Estimate for average student	16.42	12.83	15.41	12.10
Predicted future number correct	17.67	14.92	16.58	12.40
Rated ability at task	4.87	5.33	6.00	4.47

Note. Twelve actor–observer pairs served in each condition except the failure–debriefing condition, in which 10 participated.

average student in the success condition ($M = 1.75$) than in the failure condition ($M = .23$), and this difference was not significantly influenced by the briefing versus debriefing manipulation.

Turning to the predicted future number correct, we again find reliable effects of feedback irrespective of the briefing versus debriefing variation. The main effect of feedback was significant, $F(1, 42) = 14.35$, $p < .001$, and no other effects or interactions approached significance. Simple main effect analyses revealed that the feedback effect was significant for both briefing, $F(1, 42) = 6.42$, $p < .05$, and debriefing, $F(1, 42) = 8.49$, $p < .01$. We also investigated the degree to which subjects' estimates of the actor's future performance departed from the baseline of their estimates of the average student's performance. Using a difference score parallel to the one noted earlier, we found no significant main effects or interactions. This failure of future estimates to depart significantly from estimates for the average student was also observed by Ross et al. (1975) in the (outcome) debriefing condition of their experiment. This is to be expected, however, because this variable was not affected by the feedback manipulation even when subjects in their experiment received no debriefing at all.

Considered together, these results indicate that briefing and debriefing have parallel influences on subjects' impressions formed from false performance feedback. In both cases, the feedback continues to inform subsequent impressions. It is worth noting that the observed success and failure means in both the briefing and the debriefing condition were within 1.5 units (estimated trials) of the comparable means in the Ross et al. outcome debriefing condition. The means for actors and observers were also similar to those of the previous study, indicating no significant main effects or interactions. And as one would expect, the effects of feedback in both the present briefing and debriefing conditions were less substantial than the effects Ross et al. observed under conditions of no debriefing.

The ability attribution measure, however, exhibited a different pattern. Ratings of the actor's ability on the judgment task did not conform to a straightforward main effect of performance feedback, as this effect was nonsignificant, $F(1, 42) = 1.86$, $p > .15$. Instead, the interaction of feedback and the timing of the denial was significant, $F(1, 42) = 7.60$, $p < .01$. For subjects in the debriefing condition, success feedback yielded greater attribution of ability to the actor ($M = 5.46$) than did failure feedback ($M = 4.41$), simple main effect $F(1, 42) = 8.91$, $p < .01$. However, for subjects in the briefing condition, this comparison was not significant; indeed, briefed subjects showed an unreliable tendency to attribute greater ability to the actor given failure feedback ($M = 5.00$) than the actor given success feedback ($M = 4.65$).

These ability findings provide evidence for

our supposition that the attributional processing of an impression would be precluded by the briefing. The ability ratings, unlike the estimates of success, call for an attributional judgment—a dispositional attribution to the actor. The absence of ability attributions consistent with the feedback among briefed subjects indicates that they stopped short of attributional processing of the feedback information; debriefed subjects, in turn, apparently engaged in just such attributional processing. So, just as Ross et al. (1975, 1977) would have it, subjects in the debriefing condition arrived at attributional judgments that were not subsequently discredited by the denial of the evidence on which those judgments were founded.

The attributional model would predict, however, that perseverancelike effects should occur only when such attributional processing of the feedback is possible. And clearly, although such processing was preempted in the briefing conditions of this experiment, subjects here nevertheless reached impressions based on the false feedback. Forewarned in the same way that debriefing subjects were later notified, subjects in the briefing condition failed to make attributions consistent with the feedback—yet based their estimates of the actor's current and future performance on evidence they were told was false.

Discussion

The present findings do not rule out the possibility that attributional processing of information might lead to increments in belief. What our findings indicate is that such processing is not necessary for the occurrence of the impression perseverance effect. The effect seems to happen because people make judgments based on denied evidence, and they do so even when the evidence is denied before it is encountered. In this sense, it is unclear whether the term *perseverance* is appropriate for the observed phenomenon. Impressions do not only *persevere,* to have a continuing influence after debriefing. They also *penetrate,* to have an intrusive influence despite a forewarning that they are false.

There are limitations on this conclusion, of course, both at the empirical and theoretical levels. Empirically, the present study must be counted as but a first step toward the specification of the nature of denial transparency. To be fair to the attributional processing model, it is necessary to note that subjects in both the debriefing and briefing conditions of this experiment may have been making certain attributional inferences that were not measured in the research. Thus, although ability attributions to the actor were clearly short-circuited by the briefing, it is conceivable that other inferences regarding the difficulty of the judgment task, the actor's mood, the degree of luck involved, or the like, may not have been similarly suspended. If this were the case, attributional processing could still be used to explain the briefing effect. Of course, Ross et al. (1975, 1977) chose ability attributions as the central element of their model because such stable, internal attributions to an actor seem the most likely (attributional) mediators of persevering impressions. For this reason, the present results invalidate the specific model they proposed and cast considerable doubt on alternative formulations that incorporate less plausible attributional mediators.

The present results do not provide an especially fine-grained view of the cognitive processes involved in denial transparency. It could be argued, for instance, that the effects of briefing and debriefing are only superficially similar, and that although debriefing effects accrue from attributional processing, briefing effects are the result of yet another process as yet unspecified. Certainly, research on the cognitive mechanisms underlying these effects is in order. Until such research is forthcoming, however, considerations of parsimony suggest that briefing and debriefing effects be considered part of the larger phenomenon of denial transparency.

Beyond these empirical limitations, the present results also suggest much unfinished theoretical work. If impressions can penetrate denials no matter when the denials are made, the observed phenomenon appears potentially to be more misleading to the perceiver than mere perseverance. For this reason, it is essential that a theoretical analysis be made of the boundary conditions of the effect. When will denial transparency lead the perceiver astray? One suggestion made earlier in this report may be worthwhile as a point of departure: When denied information is new to the recipient—in that other domain-relevant information is not available in mem-

ory—the denied information seems to have special power. This is particularly true when subsequent events make the denied information the only useful information for judgment.

The standard perseverance paradigm fits this theoretical template precisely. Subjects in such a study are asked to answer questions for which the denied information is the most clearly relevant information available. It seems safe to say that most subjects probably have not made judgments of performance on a suicide note discrimination task in their lifetimes. Their recent experience thus serves as an important anchor (Kahneman & Tversky, 1973) for subsequent judgments. No doubt, subjects refer to this experience repeatedly as they attempt to answer the various questions they are posed. And although they may well remember to attach denials to certain recollections, it seems there are many inferences from their recent experience that may slip by undenied. The denial transparency model suggests that subjects might be properly disabused of such information only if they were given counterinformation or were encouraged to reconceptualize the information and its denial in terms of a more encompassing idea that would reverse the information.

Such measures could prove useful in reducing the impact of the deceptive experimental manipulations that served as the original impetus for debriefing research (Walster et al., 1967). The present findings contribute to this original concern in yet another way, however, for they suggest that the ethical questions surrounding deception may need to be rethought. Deceptive researchers are easily portrayed as villains, after all, when the deceptive feedback they give to subjects is revealed, after the fact, to have been a lie. This is the case in the postexperimental debriefings examined by Walster et al. (1967). But is it lying to tell subjects in advance that their experimental feedback will be false? When subjects come to accept deceptive feedback despite a briefing, the experimenter's culpability for their condition becomes more difficult to estimate. Experimenters who deceive subjects may not be at fault for lying as much as they are responsible for providing subjects with information they might not have otherwise encountered. The nature of the problem of deceptive research may need to be reconsidered with this in mind.

In the end, perhaps it is not surprising that subjects who obtain unprecedented information use it even though it is known to be false. Scientists also seem to grasp whatever relevant paradigm they can find, embracing it blindly despite contradictions and empirical reversals until a better paradigm becomes available (Kuhn, 1962). In this light, it is ironic that impression perseverance has become known as a bias of the human information processor, a departure of the layperson from the normative standards of science. The phenomenon mirrors one of the most celebrated operating procedures of scientific inquiry: An old theory that is known to be false will not be discarded until a new theory is developed to take its place.

References

Anderson, C. A. (1983). Innoculation and counterexplanation: Debiasing techniques in the perseverence of social theories. *Social Cognition, 1,* 126–139.

Anderson, C. A., Lepper, M. R., & Ross, L. (1980). Perseverence of social theories: The role of explanation in the persistence of discredited information. *Journal of Personality and Social Psychology, 39,* 1037–1049.

Bateson, G. (1968). Redundancy and coding. In T. A. Sebeok (Ed.), *Animal communication: Techniques of study and results of research* (pp. 614–626). Bloomington: Indiana University Press.

Carroll, J. S. (1978). The effect of imagining an event on expectations for the event: An interpretation in terms of the availability heuristic. *Journal of Experimental Social Psychology, 14,* 88–96.

Fleming, J., & Arrowood, A. J. (1979). Information processing and the perseverence of discredited self-perceptions. *Personality and Social Psychology Bulletin, 5,* 201–205.

Freud, S. (1961). Negation. *Standard edition* (Vol. 19, pp. 235–239). London: Hogarth. (Original work published 1925)

Geiselman, R. E., Bjork, R. A., & Fishman, D. L. (1983). Disrupted retrieval in directed forgetting: A link with posthypnotic retrieval. *Journal of Experimental Psychology: General, 112,* 58–72.

Gregory, W. L., Cialdini, R. B., & Carpenter, K. M. (1982). Self-relevant scenarios as mediators of likelihood estimates and compliance. *Journal of Personality and Social Psychology, 43,* 89–99.

Hatvany, N., & Strack, F. (1980). The impact of a discredited key witness. *Journal of Applied Social Psychology, 10,* 490–509.

Jelalian, E., & Miller, A. G. (1984). The perseverence of beliefs: Conceptual perspectives and research developments. *Journal of Social and Clinical Psychology, 2,* 25–56.

Jennings, D. L., Lepper, M. R., & Ross, L. (1981). Persistence of impressions of personal persuasiveness: Perseverence of erroneous self-assessments outside the debriefing paradigm. *Personality and Social Psychology Bulletin, 7,* 257–263.

Johnson-Laird, P. N., & Tridgell, J. (1972). When negation is easier than affirmation. *Quarterly Journal of Experimental Psychology, 23,* 87–91.

Kahneman, D., & Tversky, A. (1973). On the psychology of prediction. *Psychological Review, 80,* 237–251.

Kuhn, T. (1962). *The structure of scientific revolutions.* Chicago: University of Chicago Press.

Nisbett, R., & Ross, L. (1980). *Human inference: Strategies and shortcomings of social judgment.* Englewood Cliffs, NJ: Prentice-Hall.

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

Ross, L., Lepper, M. R., Strack, F., & Steinmetz, J. (1977). Social explanation and social expectation: Effects of real and hypothetical explanations on subjective likelihood. *Journal of Personality and Social Psychology, 35,* 817–829.

Schneidman, E. S., & Farberow, N. L. (1957). *Clues to suicide.* New York: McGraw-Hill.

Semin, G. R., & Rosch, E. (1981). Activation of bipolar prototypes in attribute inferences. *Journal of Experimental Social Psychology, 17,* 472–484.

Sherman, S. J., Skov, R. B., Hervitz, E. F., & Stock, C. B. (1981). The effects of explaining hypothetical future events: From possibility to probability to actuality and beyond. *Journal of Experimental Social Psychology, 17,* 142–158.

Snyder, M., & Swann, W. B., Jr. (1978). Hypothesis testing processes in social interaction. *Journal of Personality and Social Psychology, 36,* 1202–1212.

Snyder, M., & White, P. (1981). Testing hypotheses about other people: Strategies of verification and falsification. *Personality and Social Psychology Bulletin, 7,* 39–43.

Smith, E. R., & Miller, F. D. (1983). Mediation among attributional inferences and comprehension processes: Initial findings and a general method. *Journal of Personality and Social Psychology, 44,* 492–505.

Tennen, H., & Gillen, R. (1979). The effect of debriefing on laboratory induced helplessness: An attributional analysis. *Journal of Personality, 47,* 629–642.

Thompson, W. C., Fong, G. T., & Rosenhan, D. L. (1981). Inadmissible evidence and juror verdicts. *Journal of Personality and Social Psychology, 40,* 453–463.

Tversky, A., & Kahneman, D. (1973). Availability: A heuristic for judging frequency and probability. *Cognitive Psychology, 5,* 207–232.

Walster, E., Berscheid, E., Abrahams, D., & Aronson, V. (1967). Effectiveness of debriefing following deception experiments. *Journal of Personality and Social Psychology, 6,* 371–380.

Wason, P. C. (1965). The contexts of plausible denial. *Journal of Verbal Learning and Verbal Behavior, 4,* 7–11.

Wason, P. C., & Johnson-Laird, P. N. (1972). *The psychology of reasoning.* Cambridge, MA: Harvard University Press.

Wegner, D. M., Wenzlaff, R., Kerker, R. M., & Beattie, A. E. (1981). Incrimination through innuendo: Can media questions become public answers? *Journal of Personality and Social Psychology, 40,* 822–832.

Wegner, D. M. (1984). Innuendo and damage to reputations. In T. C. Kinnear (Ed.), *Advances in Consumer Research* (Vol. 11, pp. 691–696). Provo, UT: Association for Consumer Research.

Received March 21, 1984
Revision received March 27, 1985 ■