

INOCULATION AND COUNTEREXPLANATION: DEBIASING TECHNIQUES IN THE PERSEVERANCE OF SOCIAL THEORIES

CRAIG A. ANDERSON
Rice University

Our social theories are often unresponsive to logical and empirical challenges. Earlier research has demonstrated that the process of creating causal explanations or general scenarios to explain observed events contributes to such unwarranted theory perseverance. An analysis of possible cognitive mechanisms underlying theory perseverance suggests that explanation processes might be used in "debiasing" techniques. It is predicted that theory perseverance would be reduced by inducing subjects to create causal explanations for both possible relationships between two social variables. College student subjects were given two case studies suggestive of either a positive or a negative relationship between preference for risk and success as a firefighter. Some subjects had only to consider the relationship suggested by their case history data. Others were induced to consider both possible relationships. Subjects forced to consider both relationships showed significantly less theory perseverance, supporting the analysis of the cognitive mechanisms underlying theory perseverance and suggesting possible applications in real-world contexts. The probable locus of the effects of the debiasing techniques and possible boundary conditions are discussed.

People's beliefs about themselves, about other people, and about relationships between events in the world around them are amazingly resilient. Such beliefs are often unaffected (or only slightly affected) by logical or evidential challenges that normatively should weaken or destroy them. A number of researchers have begun investigating this phenomenon of "belief perseverance" in several domains. (See Ross & Anderson, 1982, for a review of much of this research.) Ross, Lepper, and Hubbard (1975) have

I would like to thank Mark Lepper and Lee Ross for their comments in the design phase of this research. Thanks are also due to William Howell and Michael Watkins for their comments on an earlier version of this paper. This research was supported in part by NSF Grant BNS-78-01211, NIMH Grant MH-36093, and by a Rice University Faculty Research Grant. Requests for reprints should be sent to Craig A. Anderson, Department of Psychology, Rice University, Houston, Tex. 77251.

demonstrated that manipulated self-perceptions of ability were virtually unaffected by the revelation that initial success or failure at the task had been manipulated by the experimenter. That is, success subjects continued to believe that they had the ability to do well, and failure subjects continued to believe that they would do poorly at the task, despite knowing that their initial outcomes had been rigged. Observer subjects in this study showed a similar perseverance effect in their ratings of actor subjects' abilities.

More recently, the belief perseverance phenomenon has been demonstrated in the domain of social theories—our beliefs about the causal relationships between social variables. Anderson, Lepper, and Ross (1980) presented subjects with purportedly authentic case histories of two firefighters. Each case history included information about the firefighter's preference for risk, indicated by his responses to a paper-and-pencil Risky-Conservative Choice Test, and information about his job performance as a firefighter. By manipulating this case history information, some subjects were led to believe in a positive relationship between risk preference and firefighting ability (high risk preference associated with being a successful firefighter, and low risk preference associated with being a failure as a firefighter). Other subjects were similarly led to believe that the relationship was a negative one. All subjects wrote an explanation of the relationship they discovered in the case histories. Half of the subjects in each condition were thoroughly debriefed concerning the fictitious nature of the case histories. At this point the debriefed subjects all had logically equivalent information about the relationship between risk preference and firefighting performance—specifically, no valid information at all. However, subsequent assessments of their personal beliefs revealed widely divergent theories. Subjects initially led to believe that a positive relationship existed continued to hold that belief, while subjects led to believe in a negative relationship clung to their initial belief as well.

Internal analyses of the Anderson *et al.* (1980) study revealed that subjects whose explanations were of a causal nature or who made reference to general causal scenarios displayed more perseverance in their initial theories than did subjects whose explanations were simply restatements of the case history data. A subsequent study (Anderson *et al.*, 1980, Experiment 2), in which causal explanation of a theory was experimentally manipulated, confirmed this finding. Creating a general causal scenario or explanation for weak data leads to a stronger tendency to cling to a discredited and normatively unwarranted theory. (See also Ross, Lepper, Strack, & Steinmetz, 1977, for an examination of explanation effects on social perception.)

But as Anderson *et al.* pointed out, many important decisions must be made on the basis of weak or contradictory data (*cf.* Abelson, 1976; Chapman & Chapman, 1969; Janis & Mann, 1977). Our typical response to

such demands is to engage in causal processing (Kelley, 1967, 1973; Pyszczynski & Greenberg, 1981), and to use the now-resistant beliefs thus generated to make these important decisions. Indeed, recent research (Anderson, in press) indicates that the type of data most likely to prove false—weak but vivid case histories—is also most likely to induce general causal processing and the most belief perseverance. Understanding the cognitive mechanisms underlying theory perseverance thus has important practical implications. By understanding the mechanisms we may be able to devise “debiasing” techniques for use in important decision-making contexts. Conversely, we may discover more about the perseverance-producing mechanisms by restructuring the tasks and observing the conditions under which perseverance is reduced or eliminated, or maximized.

As a starting point for constructing potential debiasing techniques, consider a possible mechanism for the perseverance-enhancing effect of explanation. It appears that general, causal explanations (or scenarios) become relatively independent of the data that gave rise to them. In this way, the theory is isolated from the data and from challenges to the data. For example, in the risk-preference/firefighter studies (Anderson *et al.*, 1980), a common explanation generated by subjects exposed to positive-relationship data was “Firefighting is a risky occupation; one must take risks to be a good firefighter. Therefore, people who prefer risk will be better firefighters.” The discovery that the case history data are fictitious does not reduce the reasonableness or apparent validity of the explanation. The “theory” thus remains as an available (Carroll, 1978; Tversky & Kahneman, 1973) and reasonable way of understanding how the various factors could be related. That is, the theory remains as a salient heuristic for making judgments.

One debiasing technique, then, might be to avoid creating causal explanations or theories. Such a suggestion is as undesirable as it is impossible. Many of our theories are quite useful, both as information organizers and as predictive tools. The problem lies not in our propensity to create theories, but in our underestimation of how easy it is to create plausible theories for any particular set of events we wish to explain. In the firefighter studies cited earlier, subjects erred not in creating explanations for the initial case history data but in failing to realize, after the debriefing, that had they been given an opposite set of bogus case history data, they would have been able to create a plausible opposite explanation or theory. Perhaps their beliefs remained relatively unchanged because they had assumed their theory was the only plausible one.

This analysis, if accurate, suggests that theory perseverance may be reduced by making salient the plausibility of alternative relationships and explanations of alternative relationships between the same set of social variables. One might induce the theory holder to create conceptually op-

posite theories in the domain of interest. For example, subjects who initially explained how risk preference and firefighting ability are positively related might show less theory perseverance if induced to explain how the variables might be negatively related. (See Koriat, Lichtenstein, & Fischhoff, 1980; Slovic & Fischhoff, 1977, for similar procedures in different contexts.)

There are two cases in which such a perseverance-reducing procedure might be introduced. One case occurs when a person has already created a theory for data that have subsequently been discredited (either wholly, as in the firefighter studies, or partially, as would be true in more real-world situations). In this case, the theory holder might be induced to imagine that there are valid data implying a conceptually opposite relationship, and to try to explain how or why such a relationship might exist. That is, the person could create a “counterexplanation.”

The second case requires a bit more foresight on the part of the decision maker. If one is about to examine information that is weak or likely to be contradicted at some later time, one may try to create plausible explanations for both (or all) possible relationships between the relevant variables prior to examining any data at all. This “inoculation” procedure may thus reduce unwarranted theory perseverance by showing how easily either relationship might be explained, and how hard it is to predict which is true.

Both of these procedures should theoretically make the decision maker aware that competing, salient, plausible theories can easily be created. This should allow the decision maker to be more flexible when initial data (and beliefs based on them) are challenged. The present study was designed to test this prediction for both types of manipulation.

METHOD

OVERVIEW

All subjects were given case history data that suggested that a positive or a negative relationship exists between a firefighter trainee's preference for risky versus conservative choices and his subsequent success as a firefighter. All subjects were later debriefed about the totally fictitious nature of these case histories. Within each of these conditions, subjects were randomly assigned to one of four explanation conditions. “No-Explanation” subjects did not write any explanations at all. “Explanation” subjects wrote only an explanation of the relationship discovered in their case history data, prior to debriefing. “Counterexplanation” subjects were identical to Explanation subjects, except that, subsequent to the debriefing, they

also wrote an explanation of the conceptually opposite relationship. "Inoculation" subjects wrote explanations of both a positive and a negative relationship prior to seeing the case histories. Finally, all subjects completed several dependent measures assessing their beliefs concerning the true relationship between these two variables and the predictive power of the relationship. The final design was thus a 2×4 factorial—relationship (positive vs. negative) \times explanation (No Explanation vs. Explanation vs. Counterexplanation vs. Inoculation).

SUBJECTS

In all, 88 undergraduates participated in the experiment in three group sessions, and received credit toward a course requirement. Subjects were randomly assigned to conditions in blocks of eight.

PROCEDURE

Subjects were told that the experiment concerned how well people are able to discover relationships between personal characteristics of people and their behavioral outcomes. They were informed that their main task was to examine some case history data and to see if they could discover any underlying relationships between general traits and specific behaviors. After questions about these general instructions were answered, the subjects were given booklets containing the experimental material.

Manipulation of Initial Theories

Instructions on each booklet informed subjects that their task was to consider the relationship between eventual success or failure as a firefighter and prior performance on a Risky-Conservative Choice Test (RCC Test). The manipulation of subjects' initial theories took place in the form of two case histories. In addition to nondiagnostic background information, the success or failure of the target firefighter and five of his "most representative" RCC Test items and responses were presented in detail. By manipulating the success or failure of each firefighter (as assessed by superiors' ratings) and the riskiness of each firefighter's RCC Test responses, subjects were presented with data that appeared to support either a positive or a negative relationship between risk preference and subsequent success as a firefighter. In positive-relationship conditions, for example, the successful firefighter made risky choices on four of his five representative RCC Test items; conversely, the failure firefighter made only one risky choice. (See

Anderson *et al.*, 1980, for a more complete description of this belief manipulation.)

After examining this information, each subject completed a rating scale designed to assess the relationship that had been "discovered in the two case studies." This scale served as a manipulation check.

Manipulation of Explanation Processes

In the No-Explanation conditions, subjects were never asked to write out any explanations of possible relationships between the two relevant variables. Of course, the process of trying to discover an underlying relationship probably guarantees that at least some causal processing will take place. In the Explanation conditions, subjects were explicitly instructed to write out an explanation of the relationship they discovered in the case history data. In the Inoculation conditions, subjects were asked to write out plausible explanations for both a positive and a negative relationship *before* looking at the case history data. In all three of these sets of conditions, subjects were then debriefed (discussed below) and their beliefs concerning the true relationship were assessed. In the Counterexplanation conditions, though, subjects were asked to explain their discovered relationship after examining the case histories, were debriefed, then wrote an explanation of how the opposite relationship might be true. This was followed by the dependent measures of beliefs. Table 1 presents the sequence

TABLE 1
Sequence of Critical Events in the Experiment

EXPLANATION CONDITION	SEQUENCE OF EVENTS				
	FIRST	SECOND	THIRD	FOURTH	FIFTH
No explanation	Examine case histories	Debriefing	Dependent measures		
Explanation	Examine case histories	Explain case history data	Debriefing	Dependent measures	
Counterexplanation	Examine case histories	Explain case history data	Debriefing	Explain opposite	Dependent measures
Inoculation	Explain both possible relationships	Examine case histories	Debriefing	Dependent measures	

of critical events for each of the four levels of the Explanation independent variable.

Debriefing and Dependent Variables

The debriefing consisted of informing subjects that (1) the case histories were fictitious, (2) both the ability ratings and the risk-preference information had been manufactured to create positive and negative relationship conditions, (3) subjects had been randomly assigned to one of these conditions, and (4) the purpose of this deception was to maximize their performance by having them perform a task with purportedly authentic data. Three measures of subjects' beliefs were obtained: a measure of the perceived criterion validity of the riskiness test; a measure of the extent to which subjects were willing to generalize their beliefs to new target cases; and a measure of the extent to which subjects were willing to generalize their beliefs to new RCC Test items. For the first measure subjects estimated the average percentage of risky choices for two groups of firefighters—those who had subsequently become successful and those who had become failures at the job. The perceived criterion validity of the RCC Test was assessed by subtracting the expected percentage of risky responses for failure firefighters from the corresponding percentage for successful firefighters. For the second measure subjects were presented with information on four new firefighter trainees, including their responses to one RCC Test item. Subjects then predicted each trainee's subsequent success or failure. These predictions were scored to assess the degree to which they would follow from a positive or negative theory. For the final measure, subjects were presented with five new RCC Test items, and were asked to predict the percentage of risky choices superior and inferior firefighters would make. A difference score served as each subject's index of theory on this measure. The debriefing and the dependent variables are more thoroughly described in Anderson *et al.* (1980).

RESULTS AND DISCUSSION

MANIPULATION CHECK

To ascertain that subjects in the positive- and negative-relationship conditions had indeed "discovered" these different relationships, a manipulation check was included immediately after the two case histories. The results from this measure—a 101-point scale anchored at Highly Positive Relationship (50), No Relationship (0), and Highly Negative Relationship (-50)—indicated that subjects in the positive-relationship conditions dis-

covered a positive relationship, $M = 29.05$, $t(80) = 9.93$, $p < .0001$,¹ while subjects in the negative-relationship conditions discovered a negative relationship, $M = -22.41$, $t(80) = 7.66$, $p < .0001$.

DEPENDENT VARIABLES

Since the three dependent measures proved to be highly correlated (average $r = .70$), the data on each were transformed into z scores and summed to provide a composite measure of subjects' beliefs concerning the true relationship between risk preference and subsequent success as a firefighter. It is on these data that the primary analyses were performed.²

In order to test the notion that theory perseverance may be reduced by inducing subjects to consider and explain both possible relationships, we must first demonstrate that the basic perseverance effect has occurred. Thus let us examine first the conditions in which subjects were led to consider only the relationship suggested by their case history data—the No-Explanation and the Explanation conditions. As shown in Figure 1, these subjects' initial theories survived the total discrediting of the evidential basis of those theories despite being based on objectively weak data (the two cases). That is, when subjects considered one side only, those initially exposed to a positive relationship continued to believe that risky responses are highly diagnostic of later success, while those initially exposed to a negative relationship believed that conservative responses predict success. This difference between final beliefs of positive- and negative-relationship groups was highly significant, $F(1, 80) = 32.30$, $p < .0001$, replicating the basic perseverance of Anderson *et al.* (1980).³

Given that a strong perseverance effect was obtained, it is now possible to test whether the Inoculation and Counterexplanation procedures

1. Except where indicated, all p levels are for two-tailed tests of significance.
2. Separate analyses parallel to those to be reported below were performed on each of the three component measures. The results from these separate analyses were generally the same as for the combined analysis, though somewhat weaker. Since examination of these individual analyses adds little to our understanding of the results, they are not reported further.
3. In this paper, the perseverance effect was stronger in Explanation than in No-Explanation conditions. The present study failed to replicate this effect, $F(1, 80) < 1$. It is not clear why this occurred. Several possibilities include: (1) slight changes in instructions that may have increased explanation processes in No-Explanation subjects or decreased such processes in Explanation subjects; (2) less consistency in understanding of and compliance with instructions in the present study due to being run in large group sessions rather than small ones. This second possibility also suggests that the error variance on the component measures might be larger in the present study than in the previous ones. A comparison of the variances confirms this speculation for all three component measures, $p < .01$, $p < .01$, and $p < .25$, respectively, for the statistical comparisons.

were successful in reducing the amount of theory perseverance by forcing subjects to consider both possible relationships. Figure 1 presents the relevant comparison means. As can be seen, subjects who only had to consider one possible relationship (the Explanation and No-Explanation conditions) showed significantly more theory perseverance than did subjects induced to consider both possible relationships (the Counterexplanation and Inoculation conditions), $F(1, 80) = 5.06, p < .03$. Specific contrasts that individually tested the efficacy of the Counterexplanation and the Inoculation procedures showed that both were effective in reducing the amount of theory perseverance; $t(80) = 1.95, p < .03$, for Counterexplanation; $t(80) = 1.73, p < .05$, for Inoculation, one-tailed.⁴

Interestingly, there was no difference between the effectiveness of the two perseverance-reducing procedures, $F(1, 80) < 1$. But as is apparent from Figure 1, these procedures did not entirely eliminate unwarranted theory perseverance. Both positive- and negative-relationship subjects continued to believe in their initial theories. That is, Counterexplanation and Inoculation subjects did tend to cling to their initial theories induced by the case histories, $F(1, 80) = 6.26, p < .01$.

CONCLUSIONS

Four general conclusions can be drawn from these results. First, inducing people to create causal explanations of opposite social theories produces more flexible and appropriate responses to challenges to those theories. Subjects who explained both a positive and a negative relationship between riskiness and firefighter success displayed significantly less reluctance to abandon their theories when the data base was shown to be fictitious than did subjects who only considered the relationship suggested by their case history data. This is the first clear demonstration that a belief perseverance bias can be reduced by restructuring the task.

Second, the success of these debiasing techniques lends support to the view that perseverance biases are based on the relative availability of plausible causal explanations or scenarios. By making competing scenarios available, the bias was significantly reduced. But must this competing theory be *opposite* in order to reduce the bias? Or will making any alterna-

4. The success of these debiasing techniques could have occurred in either of two ways. Considering both possible relationships could have led to modest reductions in the perseverance bias for most subjects, or to major reductions in the bias for only a few subjects. The latter possibility would be revealed by large variance estimates for the Counterexplanation and Inoculation conditions, relative to the Explanation and No-Explanation conditions. A comparison of these variances failed to yield a significant difference, $F(40, 40) = 1.52, p > .05$, suggesting that the modest debiasing effects occurred for most of the subjects.

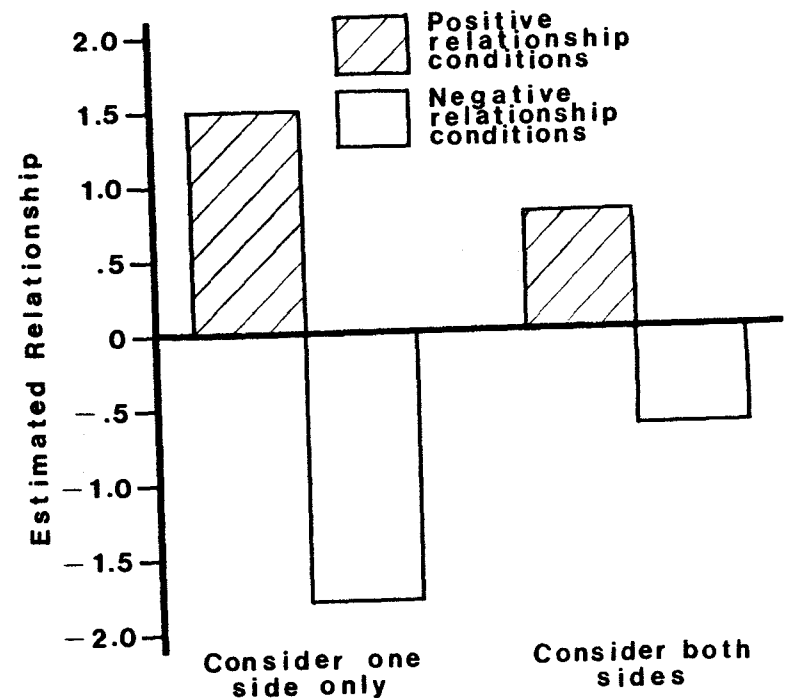


FIGURE 1

Mean Composite Index of Subjects' Personal Estimates of the True Relationship between Risk Preference and Firefighter Success. Positive Scores Indicate a Belief in a Relatively More Positive Relationship; Negative Scores Indicate a Belief in a Relatively More Negative Relationship.

tive theory available (via explanation) reduce the initial bias? To the extent that final beliefs depend upon the relative availability of various causal theories, we should expect that increasing the availability of any alternative theory, whether it is opposite to or simply different from the initial one, should lead to a decrease in the strength of the initial theory. Thus, inducing subjects to explain how a preference for physical work leads to high success as a firefighter should be an effective debiasing manipulation in this paradigm. Note, however, the caveat implicit throughout this paper that now must become explicit. The apparent plausibility of a causal theory must be considered. A ridiculous theory will most likely have little impact, regardless of its availability. (But see Hasher, Goldstein, & Toppino, 1977, for research suggesting that presentation frequency and availability are positively related to perceived truthfulness.) The "physical-work" manipulation suggested above may be an effective one only if it is seen as a plausible causal theory. The opposite theories of the present

study were effective as alternative explanations because they were plausible and made more available.

Third, the effectiveness of these procedures does not appear to depend upon when the competing explanations are generated. But while the Counterexplanation and the Inoculation procedures did not differ in overall effectiveness, they may operate in slightly different ways. The former procedure comes into play only after the initial theory has been formed and challenged. Its effectiveness may derive from the close contiguity of the challenge to the data base and the realization (via the counterexplanation) that an opposite causal relationship is also plausible. The latter procedure, however, may derive its effectiveness at the theory formation stage as well. That is, the inoculating explanations may improve people's foresight (Fischhoff, 1975), allowing them to avoid being overly swayed by incoming data that are weak and, hence, likely to be contradicted. If this speculation is correct, we should expect the "discovered" theories of the Inoculation subjects to be less extreme at the manipulation check. This does in fact appear to be the case in the present study. While the mean difference between the positive- and negative-relationship groups for the Explanation subjects is 62.64, the corresponding mean difference for Inoculation subjects is only 38.81, $t(80) = 2.04$, $p < .05$. The Inoculation procedure appears to make people somewhat more cautious in forming theories on the basis of weak data.

Thus, a fourth conclusion is that while the Counterexplanation procedure has its impact on theory revision, the Inoculation procedure has at least part of its impact on initial theory formation.

The applicability of these results to real-world decision contexts depends upon a number of factors and potential boundary conditions. One issue, for instance, is whether these procedures continue to be as effective when the theory involved has a strong emotional component, such as beliefs concerning the deterrent effect of capital punishment. Such emotional commitment may prevent one from considering competing theories even when explicitly asked to do so. Certain situations are also likely to elicit one-sided thinking as a result of emotional commitment. In many decision-making groups there may be strong pressure on participants to think and act in the same way, preserving the harmonious nature of the group as exemplified by the "groupthink" phenomenon (Janis, 1972).

On the other hand, some situations demand that the decision maker consider all alternatives. For example, arbitrators are expected to be fair, and thus, to consider equally the disputes and explanations of all sides. In such situations perseverance biases may be less serious, as "counterexplanation" procedures are sometimes a part of the system.

Simply demanding that one be accurate or fair, however, does not guarantee that all points of view will be considered. Lord, Lepper, and

Thompson (1980), for instance, found that instructions to be "as objective and unbiased as possible" in evaluating new studies related to capital punishment failed to reduce biased assimilation processes. Indeed, in that experiment biased evaluation of new studies was eliminated only when subjects were instructed to "ask yourself at each step whether you would have made the same . . . evaluations had exactly the same study produced results on the other side of the issue." In other words, the assimilation bias was reduced only when subjects were explicitly instructed to consider the new evidence from both points of view. Thus, while emotional commitment may impede the use of debiasing strategies, it does not preclude the use of such strategies.

A second issue concerns the time frame in which the various events occur. If the empirical evidence for a theory is challenged months after initial theory formation, will a counterexplanation procedure have any impact? To the extent that the theory has become assimilated with other knowledge structures of the individual, we might expect little impact of either the empirical challenge or the counterexplanation. An empirical answer, however, awaits further research.

At a more theoretical level, there are a host of very difficult questions concerning the normativeness of belief change (or perseverance) under varying conditions. One could argue in the present study, for instance, that both the theory perseverance and the reduction in theory perseverance as a result of creating opposite explanations are counternormative effects. Adopting a theory on the basis of only two dubiously representative case histories is normatively suspect. Clinging to that theory when the case histories are shown to be worthless is clearly counternormative. Reducing one's belief in that theory because of an explicit request to consider how the opposite might be true is also counternormative in the sense that no new data have been added. More precise specifications of what is normatively appropriate are needed, though. (See Ross & Lepper, in press, for a discussion of many of these considerations.)

A final, intriguing issue concerns the applicability of these explanation procedures to self-perceptions. Can the perseverant, negative self-assessments of psychology subjects in "failure" conditions (Ross, Lepper, & Hubbard, 1975) be reduced by having them consider why they might actually be very skilled at the task? While there appear to be no studies that directly address this question, a host of attribution studies have shown that inducing subjects to consider alternative reasons for their successes and failures leads to different reactions to these outcomes (e.g., Anderson & Jennings, 1980; Chapin & Dyck, 1976; Diener & Dweck, 1978). In general, then, we might expect that increasing the availability of plausible alternative explanations for an event will reduce both the strength of the initial belief about the event and the consequences of that belief.

REFERENCES

- Abelson, R. P. Script processing in attitude formation and decision-making. In J. S. Carroll & J. W. Payne (Eds.), *Cognitive and social behavior*. Hillsdale, N.J.: Erlbaum, 1976.
- Anderson, C. A. Abstract and concrete data in the perseverance of social theories: When weak data lead to unshakeable beliefs. *Journal of Experimental Social Psychology*, in press.
- Anderson, C. A., & Jennings, D. L. When experiences of failure promote expectations of success: The impact of attributing failure to ineffective strategies. *Journal of Personality*, 1980, 48, 393-407.
- Anderson, C. A., Lepper, M. R., & Ross, L. Perseverance of social theories: The role of explanation in the persistence of discredited information. *Journal of Personality and Social Psychology*, 1980, 39, 1037-1049.
- Carroll, J. S. The effect of imagining an event on expectations for the event: An interpretation in terms of the availability heuristic. *Journal of Experimental Social Psychology*, 1978, 14, 88-96.
- Chapin, M., & Dyck, D. G. Persistence in children's reading behavior as a function of N length and attribution retraining. *Journal of Abnormal Psychology*, 1976, 85, 511-515.
- Chapman, L. J., & Chapman, J. P. Illusory correlation as an obstacle to the use of valid psychodiagnostic signs. *Journal of Abnormal Psychology*, 1969, 74, 271-280.
- Diener, C. I., & Dweck, C. S. An analysis of learned helplessness: Continuous changes in performance, strategy, and achievement cognitions following failure. *Journal of Personality and Social Psychology*, 1978, 36, 451-462.
- Fischhoff, B. Hindsight \neq foresight: The effect of outcome knowledge on judgment under uncertainty. *Journal of Experimental Psychology: Human Perception and Performance*, 1975, 1, 288-299.
- Hasher, L., Goldstein, D., & Toppino, T. Frequency and the conference of referential validity. *Journal of Verbal Learning and Verbal Behavior*, 1977, 16, 107-112.
- Janis, I. L. *Victims of groupthink*. Boston: Houghton Mifflin, 1972.
- Janis, I. L., & Mann, L. *Decision making: A psychological analysis of conflict, choice, and commitment*. New York: Free Press, 1977.
- Lord, C. G., Lepper, M. R., & Thompson, W. C. *Inhibiting biased assimilation in the consideration of new evidence on social policy issues*. Paper presented at the Annual Convention of the American Psychological Association, Montreal, September 1980.
- Kelley, H. H. Attribution theory in social psychology. In D. Levine (Ed.), *Nebraska symposium on motivation* (Vol. 15). Lincoln: University of Nebraska Press, 1967.
- Kelley, H. H. The process of causal attribution. *American Psychologist*, 1973, 28, 107-128.
- Koriat, A., Lichtenstein, S., & Fischhoff, B. Reasons for confidence. *Journal of Experimental Psychology: Human Learning and Memory*, 1980, 6, 107-118.
- Pyszczynski, T. A., & Greenberg, J. Role of disconfirmed expectancies in the instigation of attributional processing. *Journal of Personality and Social Psychology*, 1981, 40, 31-38.
- Ross, L., & Anderson, C. Shortcomings in the attribution process: On the origins and maintenance of erroneous social assessments. In D. Kahneman, P. Slovic, & A. Tversky (Eds.), *Judgment under uncertainty: Heuristics and biases*. New York: Cambridge University Press, 1982.
- Ross, L., & Lepper, M. R. The perseverance of beliefs: Empirical and normative considerations. In R. A. Schweder & D. Fiske (Eds.), *New directions for methodology of behavioral science: Fallible judgment in behavioral research*. San Francisco: Jossey-Bass, in press.
- Ross, L., Lepper, M. R., & Hubbard, M. Perseverance in self-perception and social perception: Biased attribution processes in the debriefing paradigm. *Journal of Personality and Social Psychology*, 1975, 32, 880-892.
- Ross, L., Lepper, M. R., Strack, F., & Steinmetz, J. Social explanation and social expectation: Effects of real and hypothetical explanations on subjective likelihood. *Journal of Personality and Social Psychology*, 1977, 35, 817-829.
- Slovic, P., & Fischhoff, B. On the psychology of experimental surprises. *Journal of Experimental Psychology: Human Perception and Performance*, 1977, 3, 544-551.
- Tversky, A., & Kahneman, D. Availability: A heuristic for judging frequency and probability. *Cognitive Psychology*, 1973, 5, 207-232.