JOURNAL OF EXPERIMENTAL SOCIAL PSYCHOLOGY 19, 93–108 (1983)

Abstract and Concrete Data in the Perseverance of Social Theories: When Weak Data Lead to Unshakeable Beliefs

CRAIG A. ANDERSON

Rice University

Received September 4, 1981

The perseverance effect—the finding that people cling to their initial beliefs more strongly than appears warranted—has been demonstrated in a wide variety of settings. The existing explanation of the effect implies that beliefs based on concrete data should be more resistant to challenges than beliefs based on abstract data. The present studies compared the amount of belief perseverance when the beliefs were initially based on either abstract or concrete data. Subjects examined either two case histories (concrete data) or a statistical summary (abstract data) suggestive of either a positive or a negative relationship between fire fighter trainees' level of preference for high risk and their subsequent success as firefighters. These data sets were equated for the initial strength of beliefs they induced. Subjects were then thoroughly debriefed about the fictitious nature of their initial data. Subsequent assessments of subjects' personal beliefs about the true relationship revealed (a) significant levels of theory perseverance both immediately and 1 week later; (b) significantly more perseverance in the concrete data conditions, both immediately and 1 week later. Experiment 2 revealed that subjects frequently engage in causal processing spontaneously, especially when examining concrete data. Overall, the data suggested that memory for initial data did not contribute to the abstract/concrete effects, but that the generation of general, causal explanations did contribute to the stronger perseverance of theories in the concrete conditions.

The human propensity to cling to initial opinions, attitudes, and theories has long been recognized (e.g., Allport, 1954; Hovland, Janis, & Kelley, 1953; Luchins, 1957). More recent evidence suggests that people cling to their initial beliefs to a degree that is normatively inappropriate (cf. Ross & Anderson, 1982). Such counternormative *belief perseverance* has been demonstrated for initial beliefs about oneself (Ross, Lepper, & Hubbard, 1975), initial beliefs about another person (Ross, Lepper, &

I thank William Howell and Michael Watkins for their helpful comments on an earlier version of this paper, and Mark Zanna for his suggestions concerning Experiment 2. This research was supported in part by a Rice University Research Grant. Address reprint requests to Craig A. Anderson, Department of Psychology, Rice University, Houston, TX 77251.

0022-1031/83/020093-16$03.00/0

Hubbard, 1975, Experiment 2), and initial social theories or beliefs about the functional relationship between social variables (Anderson, Lepper, & Ross, 1980; Anderson, 1982).

For example, in the domain of social theories, Anderson et al. (1980) manipulated subjects' initial beliefs about the relationship between risk preference and job performance as a firefighter. This manipulation was accomplished by the presentation of two detailed case histories of fire fighters. Subjects' theories about the true relationship between risk preference and job performance were only slightly affected by the revelation that the case histories were fictitious. Those initially led to believe that high riskiness predicted future success continued to believe in this positive relationship; those initially led to believe that high riskiness predicted future failure continued to believe in this negative relationship. But since subjects in both conditions knew that the initial data were fictitious, the final beliefs of these groups should not have differed, according to normative models of decision making. (See Ross & Lepper, in press, for a more thorough discussion of normative considerations in belief perseverance.)

While belief perseverance may sometimes serve motivational ends, as in racial prejudice, recent evidence suggests that more purely cognitive processes may also lead people to persist in holding on to their initial beliefs. Using the risk preference/firefighter materials described earlier, Anderson et al. (1980) examined the effects of the cognitive process of creating causal explanations or scenarios on the perseverance of social theories. Some subjects were induced to write out such explanations; others were not. Subjects who had engaged in this causal processing were virtually unaffected by the revelation that the data base for their theory was totally fictitious. That is, explanation-generating subjects showed significantly more theory perseverance than did subjects who were not explicitly induced to engage in such cognitive work. Once created, the explanations became functionally independent of the data that instigated their creation, remaining as highly available (Tversky & Kahneman, 1973) cues when final beliefs were assessed (cf. Carroll, 1978). Similar explanation effects have been demonstrated with beliefs about other people (Ross, Lepper, Strack, & Steinmetz, 1977).

While unwarranted perseverance effects have been demonstrated across wide variety of types of beliefs (e.g., about oneself, about other people, about functional relationships between social variables), there are no studies of how different types of data may affect the perseverance of beliefs induced by the data. The initial belief inductions of all previous studies have been based on data that are concrete and vivid, and usually quite weak as well. For instance, the theory perseverance studies (Anderson et al., 1980) induced initial beliefs in either a positive or a negative relationship between risk preference and success as a fire fighter by presenting only two case histories—a rather small sample of dubious

representativeness. These case histories may also be seen as containing relatively concrete information, such as each person's hobbies and his responses to the "Risky–Conservative Choice Test." In real world situations data of this type are notoriously unreliable, and are quite likely to be challenged on both logical and empirical grounds. But as the research shows, data of this type also lead to beliefs that persevere in the face of such challenges.

Beliefs and decisions are sometimes based on more reliable data that are also more abstract—for instance, statistical summaries of empirical studies. Such abstract data are less likely to be challenged logically or disconfirmed empirically, so we might expect beliefs based on such reliable data to be held more strongly, or to be more susceptible to perseverance biases.

But a number of researchers in both the cognitive (D'Agostino, O'Neill, & Pavio, 1977; Parker, 1981; Smith, 1981) and the social (Borgida & Nisbett, 1977; Enzle, Hansen, & Lowe, 1975; Hamill, Wilson, & Nisbett, 1980; Reyes, Thompson, & Bower, 1980) areas have demonstrated that vivid and concrete information can have a greater impact than abstract and pallid information on memory, initial person impressions, and initial social theories. Even logically inferior information can predominate, if it is more vivid or concrete than its logically superior competition. All these effects presumedly result from people's reliance on the availability heuristic (Tversky & Kahneman, 1973); and concrete information is generally more memorable or available than abstract information. This suggests that belief perseverance may be stronger when the initial belief is based on concrete data than when based on abstract data, even if the concrete data is logically (and statistically) inferior.

Another possible difference between concrete and abstract data bases is the relative ease of creating general causal explanations or scenarios for each. Concrete data, such as case histories, contain rich detail that may make explanations easier to generate or more likely to occur. As we have seen earlier, though, creating general explanations also increases belief perseverance by isolating the belief from the data that led to the formation of the belief.

Thus, because of both the memorability and the explanation induction potential of concrete data, we are led to the rather paradoxical hypothesis that people may persevere more in their initial belief when it is based on weak but concrete data than when it is based on strong but abstract data. The present studies were conducted to examine theory perseverance phenomena that result from these two different types of data. In particular, theory perseverance in the risk preference/fire fighter performance paradigm was examined as a function of type of initial data base—vivid, concrete case histories versus pallid, abstract statistical summaries.

Four related questions were addressed by the present studies. First,

can theories that are based on abstract, statistical data survive the total discrediting of that initial data base? That is, does theory perseverance occur for abstract data?

Second, is theory perseverance stronger when the initial theory is based on inferior, but concrete, case history data than when based on superior, but abstract, statistical data? Here we are interested in the relative perseverance of theories that are *initially* of equal perceived strength. Since concrete information often has more impact than logically equivalent abstract information, it becomes necessary to present abstract data that are stronger than the concrete data to produce initial theories of equal perceived strength. Once the abstract and concrete data bases have been equated in this fashion (through pretesting), we can compare the relative amounts of theory perseverance subsequent to the discrediting of the data. The prediction is that concrete data will lead to more general causal processing and to more theory perseverance.

Third, will the amount of theory perseverance in both abstract and concrete conditions change over time? Reyes et al. (1980) reported that their vividness manipulation had its impact only on delayed measures, as a result of the differential memorability of vivid and nonvivid information. In the first study to be reported, subjects' theories were assessed both immediately after the theory was induced and the data discredited, and after 1 week had passed. Since abstract data are not very memorable and are less likely to induce general causal analysis, we might expect relatively less theory perseverance in the abstract conditions after a 1-week delay. Conversely, since the concrete data are both memorable and likely to elicit general causal analysis, we might expect the amount of theory perseverance in the concrete conditions to remain at a high level.

Fourth, the role of explicit instructions to write out one's explanations was further examined. As pointed out by Anderson et al. (1980), thinking of the relationship in causal terms may lead to more theory perseverance regardless of whether or not anything is written out. The explanation manipulation presumably increased perseverance in that study by increasing the proportion of subjects who engaged in such causal thinking. To test these notions, some subjects were explicitly asked to write out explanations, while some were not (Study 1). Finally, several measures assessed the extent to which subjects spontaneously engaged in causal thinking in the absence of explanation instructions (Study 2).

EXPERIMENT 1

Method

Overview

Subjects examined data suggesting that either a positive or a negative relationship exists between a trainee's level of risk preference and his subsequent performance as a firefighter.

This initial theory induction was presented either in the form of two concrete case histories or an abstract statistical summary of twenty cases. Approximately half of the subjects were asked to write out an explanation of the relationship they had discovered in this initial data. All subjects were then thoroughly debriefed about the totally fictitious nature of the initial data. Subsequently, they completed several prediction tasks designed to assess their beliefs concerning the true relationship between risk preference and firefighting performance. One week later these same dependent measures were again obtained.

The basic design was thus a 2 × 2 × 2 factorial—Relationship (Positive vs Negative) by Data Type (Abstract vs Concrete) by Explanation (Explanation vs No Explanation).

Subjects

Fifty male and forty-eight female undergraduates participated in the experiment in group sessions ranging in size from 8 to 24 people, and received credit toward a course requirement. Within each session, subjects were randomly assigned to the various conditions in blocks of eight. The experimenter was unaware of the subjects' conditions.

Procedure

Subjects were told that the experiment was concerned with how well people are able to "discover relationships between personal characteristics and behavioral outcomes, depending upon how strong the relationship appears to be in sample data." They were informed that their main task was to examine the sample data presented to them and to try to discover the underlying relationship between the personal characteristic and behavioral outcome specified. After elaborating on these general points and answering any questions, the experimenter gave subjects booklets containing the experimental materials.

Manipulation of initial theories. Instructions on each booklet informed subjects that their task was to examine the relationship between trainees' risk preferences, as measured by the "Risky–Conservative Choice Test" (RCC test), and eventual performance as a firefighter. The data on risk preference and success as a firefighter were varied systematically, creating data sets indicative of either a positive relationship (i.e., high risk—high performance; low risk—low performance) or a negative relationship (i.e., high risk—low performance; low risk—high performance). Immediately thereafter, a rating scale was administered to assess what relationship the subject had "discovered in the case studies." This scale provided a check on the equality of the initial theory inductions between the abstract and concrete conditions.

Manipulation of data type. Subjects in the Concrete data conditions received detailed case histories of two firefighters. In addition to nondiagnostic background information, the case histories included overall job performance ratings, and five of the target fire fighter's "most representative" RCC test items and his responses to the items. (See Anderson et al., 1980, for a more complete discussion of this manipulation.)

Subjects in the Abstract data conditions received a sheet summarizing a study of the relationship between risk preference and success as a fire fighter. The RCC test was described as a 25-item test, with each item presenting a hypothetical situation that demands a choice of either a risky or a conservative course of action (this same description was given to concrete data subjects also). Next, raw data on 20 fire fighters (some successes and some failures at the job) were presented, as were the overall average number of risky and conservative choices made by the successful and the failure fire fighters.

From a statistical standpoint, the abstract data were quite superior to the concrete data, being based on a larger sample and having equivalent risky (or conservative) means for success and failure firefighters. (Indeed, a t test on the abstract data yields highly significant differences, $t(18) = 8.10$, $p < .0001$.) The abstract data, however, lacked the rich detail and explanation-evoking potential of the concrete data.

Explanation manipulation. Subjects in the Explanation conditions were asked explicitly

to write out a general explanation of the relationship they discovered in their initial data. Subjects in the No Explanation conditions were not asked to write out explanations of their discovered relationship. Of course, the latter subjects could also generate explanations as they examined the data. Anderson et al. (1980) showed the requiring written explanations can increase theory perseverance, presumably by increasing the probability that a given subject will engage in causal processing. As a test of the hypothesis that examining concrete data leads to more causal processing than examining abstract data, the written explanations were classified as being either of a general causal type or as not being causal. This was done by a rater blind to the experimental conditions of the subjects.

Debriefing and dependent variables. Prior to completing any dependent measures, all subjects were debriefed concerning the fictitious nature of the initial data. More specifically, subjects were informed that both the job performance and the risk preference data had been manufactured by the experimenters. To further ensure that subjects did not perceive the initial data to be representative of a true relationship, they were also informed that they had been randomly assigned to a condition that provided information consistent with either a positive or a negative relationship. The rationale of this deception was purportedly to maximize subject performance at the discovery task by having them perform the task on data they believed to be authentic.

Subsequent to this debriefing, three measures of subjects' personal beliefs about the "true" relationship were obtained. The first measure was a general measure of the perceived criterion validity of the riskiness test. Subjects indicated their belief on a 7-point scale, with 1 indicating a belief in a very negative relationship, and 7 indicating a belief in a very positive relationship. The second and third measures dealt with subjects' willingness to use their theory in making predictions about five new case histories and six new RCC test items. (See Anderson et al., 1980, for a more complete description of the debriefing and the latter two dependent variables.)

When subjects returned to the lab 1 week later, their theories were assessed with these same dependent measures after being told that their current answers "need not be the same as last week—that is, it is perfectly acceptable to give different answers. Of course, the same answers may be given also." Each subject also indicated which relationship they had been presented with in their initial booklet 1 week earlier, as a memory measure.

Finally, subjects were probed for suspicion and were given a thorough explanation of the study and of the processes that may mediate the unwarranted perseverance of initial beliefs.

Note that subjects were explicitly made aware of the private nature of their responses, thus minimizing impression management concerns.

Results and Discussion

Manipulation Check

Immediately after examining the initial data, subjects indicated the direction and strength of the relationship they had discovered on a 101-point scale anchored at "Highly Positive Relationship" (+50), "No Relationship" (0), and "Highly Negative Relationship" (−50). Results from this measure indicated that subjects in Positive Relationship conditions discovered a positive relationship, $M = 33.30$, $t(90) = 12.79$, $p < .0001$, while subjects in Negative Relationship conditions discovered a negative relationship, $M = -29.42$, $t(90) = 11.77$, $p < .0001$.[1]

[1] Unless otherwise indicated in the text, all significance levels are based on two-tailed tests.

Analyses also revealed that the magnitude of discovered theories was not different for the abstract and concrete manipulations, as indicated by the nonsignificant Relationship by Data Type interaction, $F(1, 90) = 2.52$. Note that the direction of these nonsignificant differences actually worked against the major hypotheses; abstract subjects "discovered" stronger relationships. (Since equating subjects on initial beliefs by covariance procedures only slightly strengthened the results, the simpler ANOVA results will be presented.)

Dependent Variables

Session 1 theories. There were no consistent or meaningful effects of subject sex or of the Explanation manipulation. Therefore, the means presented in all tables and figures have been collapsed across these variables. Since the three dependent measures proved to be highly intercorrelated (average $r = .66$), the data on each were transformed into z scores and summed to provide a composite measure of each subjects' beliefs concerning the true relationship between risk preference and subsequent performance as a fire fighter. An unweighted means analysis of variance was then conducted on these composite scores.[2]

The first question of interest was whether subjects displayed significant amounts of theory perseverance. As can be seen by the individual measures in Table 1 and the composite measures in Fig. 1, subjects did cling to their initial theory despite the total discrediting of the evidential base of the theory, $F(1, 90) = 28.43$, $p < .0001$.

Planned contrasts revealed that theory perseverance occurred both when the initial data were concrete case histories, $t(90) = 5.43$, $p < .0001$, and when the initial data were abstract statistical summaries, $t(90) = 2.11$, $p < .05$. Regardless of the type of data that led to initial theory formation, subjects in positive relationship conditions continued to believe that high risk predicts success while subjects in negative relationship conditions continued to believe that low risk predicts success.

The second question concerns the relative amount of theory perseverance in the abstract and concrete conditions. The appropriate test is the interaction between initial data (positive vs negative) and type of data (abstract vs concrete). As predicted, subjects given the concrete data showed significantly more theory perseverance than the abstract data subjects $F(1, 90) = 5.50$, $p < .03$. It thus appears that people hold most rigidly to their initial theories in precisely those conditions that are most likely to yield challenges to the data—when the theory is based upon weak, concrete case history data.

Session 2 theories. After a 1-week delay period, subjects returned to

[2] Separate analyses on the individual measures generally yielded the same (but weaker) effects as those to be reported on the composite measure.

TABLE 1
MEAN POSTEXPERIMENTAL BELIEFS CONCERNING THE RELATIONSHIP BETWEEN RISK PREFERENCE AND FIRE FIGHTER SUCCESS: SESSION 1

	Positive relationship		Negative relationship	
Dependent measure	Abstract	Concrete	Abstract	Concrete
Perceived criterion validity[a]	5.38	5.87	4.36	4.54
Generalization to new cases[b]	.42	2.22	−.12	−1.92
Generalization to new items[c]	19.5	36.0	1.4	−6.2
n	24	23	25	26

Note. Larger scores indicate a belief in a more positive relationship.

[a] Assessed by a 7-point scale; 1 = very negative relationship, 7 = very positive relationship.

[b] (Number of success–risky + number of failure–conservative) − (number of success–conservative + number of failure–risky) predictions to five new cases. Range of possible scores is 5 to −5.

[c] Subjects' predicted percentage of risky responses on six new items for superior minus unsuccessful fire fighters. Range of possible scores is 100 to −100.

the lab and again completed the three dependent measures. The measures once again were highly intercorrelated (average $r = .67$), so a z score composite measure was computed. Examination of the individual measures in Table 2 and the composite measure in Fig. 2 reveals, first, that the perseverance main effect persisted over the delay period $F(1, 84) = 23.53$, $p < .0001$.[3] As in Session 1, the perseverance effect was stronger when the initial theory was based on the concrete case histories rather than on the abstract summaries, $F(1, 84) = 9.71$, $p < .002$. Planned contrasts revealed that while the concrete data subjects showed the perseverance effect quite strongly, $t(84) = 5.63$, $p < .0001$, the abstract data subjects showed little effect $t(84) = 1.23$, ns. These results suggest that abstract data, being less memorable and less explanation provoking than the concrete data, yield relatively less of a perseverance bias over a delay period. To test this notion further, the data were entered into a repeated measures analysis. The specific contrast testing the prediction that abstract data subjects would show relatively less perseverance over time supported this analysis, although not impressively so $t(84) = 1.46$, $p < .08$, one tailed.

Mediating variables. Surprisingly, the Explanation manipulation did not effect the amount of perseverance in either session, F's < 1. That is, contrary to previous findings, subjects who were asked to write out their explanations of the relationship showed only slightly (and nonsignificantly) more theory perseverance than those not asked to do so. One possible reason for this may have been that equal proportions of subjects in

[3] The smaller degrees of freedom results from a few subjects who failed to return to Session 2. Deleting these subjects from the Session 1 analyses does not appreciably change any of the results or conclusions.

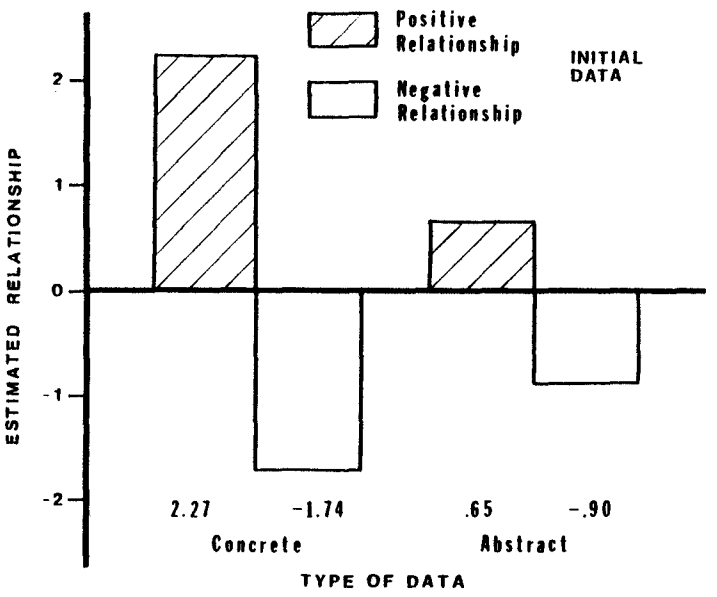

FIG. 1. Mean composite indices (z scores summed across the three measures) of subjects' personal estimates of the true relationship between risk preference and fire fighter success, Session 1.

explanation and in no explanation conditions may have engaged in causal processing. With the present data, however, we have no way to test this possibility. Experiment 2 addresses this issue.

The finding of more theory perseverance among the concrete than the abstract conditions warrants further attention. Other research suggests at least two possible causes.

TABLE 2

MEAN POSTEXPERIMENTAL BELIEFS CONCERNING THE RELATIONSHIP BETWEEN RISK PREFERENCE AND FIRE FIGHTER SUCCESS: SESSION 2

	Positive relationship		Negative relationship	
Dependent measure	Abstract	Concrete	Abstract	Concrete
Perceived criterion validity[a]	5.04	5.94	5.08	4.23
Generalization to new cases[b]	.42	2.44	−.25	−1.85
Generalization to new items[c]	17.0	29.8	.9	−5.0
n	24	18	24	26

Note. Larger scores indicate a belief in a more positive relationship.

[a] Assessed by a 7-point scale; 1 = very negative relationship, 7 = very positive relationship.

[b] (Number of success–risky + number of failure–conservative) − (number of success–conservative + number of failure–risky) predictions to five new cases. Range of possible scores is 5 to −5.

[c] Subjects' predicted percentage of risky responses on six new items for superior minus unsuccessful fire fighters. Range of possible scores is 100 to −100.

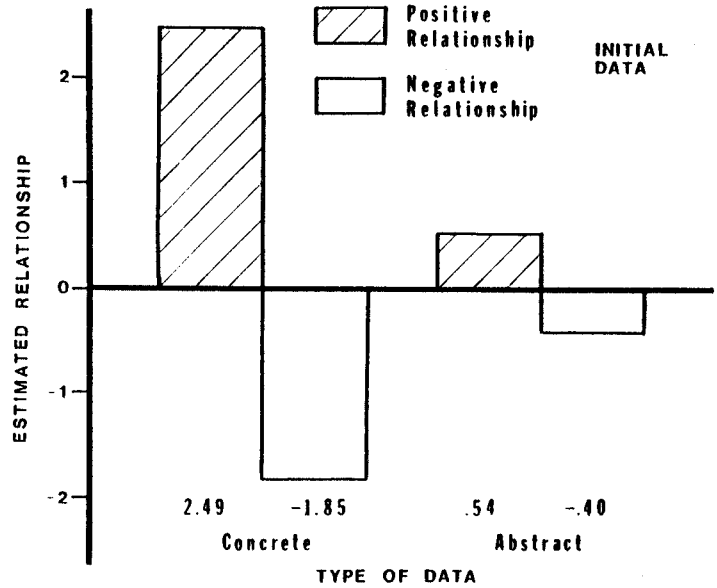

FIG. 2. Mean composite indices (z scores summed across the three measures) of subjects' personal estimates of the true relationship between risk preference and fire fighter success, Session 2.

First, concrete information may be more memorable than abstract information (e.g., D'Agostino, O'Neill, & Pavio, 1977; Parker, 1981; Reyes, Thompson, & Bower, 1980). The relative memorability of initial data may not be important in the present paradigm, though, since the memory load was not particularly high and since the total debriefing informs the subject that those data are irrelevant anyway. Second, concrete information may stimulate causal processing in the form of causal scripts or explanations (Anderson, Lepper, & Ross, 1980; Ross & Anderson, 1982; Ross, Lepper, & Hubbard, 1975), while abstract information may be less likely to do so.

While the present study was not explicitly designed to test these competing hypotheses, supplementary analyses provide some suggestive hints. Recall that at the completion of Session 2, subjects were asked to recall what relationship their original data suggested. If the memorability hypothesis is correct, then we would expect abstract data subjects to be less accurate in recalling what relationship their initial data suggested. However, abstract data subjects were not less accurate than concrete data subjects; 79 and 81% were accurate, respectively, $\chi^2(1) < 1$, ns. Thus, there was no evidence that memorability of initial data influenced the degree of perseverance shown in abstract and concrete conditions.

The causal processing hypothesis was addressed by examining the written explanations of subjects in the Explanation conditions. Recall

that each explanation was classified as either a general, causal explanation or as not a causal explanation. The explanation instructions to subjects in the abstract and concrete conditions had, of course, been identical. If the causal processing hypothesis is correct, though, abstract data subjects should find causal explanation more difficult or less clear than concrete data subjects, and hence should be less likely to write causal explanations. The data quite clearly supported this interpretation. While 81% of concrete data subjects wrote general, causal explanations, only 33% of abstract data subjects did so, $\chi^2(1) = 10.29$, $p < .001$.

In sum, Experiment 1 clearly demonstrated that theory perseverance is stronger when based upon concrete rather than abstract data, and that such perseverance can persist over a relatively long period of time (1 week). The results further suggest that the abstract/concrete difference may be due to differential causal processing. To further investigate both the effects of causal processing on theory perseverance, and the failure of the explanation manipulation in Experiment 1, a second experiment was conducted.

EXPERIMENT 2

Method

Overview

In previous research on theory perseverance, it has been assumed that subjects often spontaneously engage in causal processing while examining data for possible relationships (Anderson, 1982; Anderson et al., 1980). Furthermore, it has been postulated that such causal processing is a major determinant of theory perseverance. The present experiment was designed to examine these issues as well as those raised by Experiment 1. Each subject received either two concrete case histories or an abstract statistical summary of 20 cases, suggesting that either a positive or a negative relationship exists between a trainee's level of risk preference and his subsequent performance as a fire fighter. After examining these data, subjects indicated the direction and strength of the relationship in the data and completed measures designed to assess the amount of causal processing they did while looking at the data. Subjects were then debriefed about the fictitious nature of the initial relationship data, and completed a measure of their personal beliefs about the true relationship between risk preference and fire fighter success.

Subjects

Twenty-six male and twenty-four female undergraduates participated in group sessions that ranged in size from one to five, for course credit. Subjects were randomly assigned to the various conditions; the experimenter was blind to each subject's condition.

Procedure

The procedures were identical to those in Experiment 1, with the following exceptions. (1) None of the subjects were asked to provide written explanations. (2) Three measures of subjects' causal thinking were taken immediately after completing the manipulation check. For the first measure, subjects were asked to list all of their thoughts that they could recall having while looking over the initial data. Two judges independently coded each subject's list either as indicating some general causal processing or no causal processing.

The judges' initial classifications were the same for 46 of the 50 subjects. After discussing the four disputed cases, both judges agreed on all 50. (Note that deleting the four disputed cases from the relevant analyses does not appreciably alter the results.) The second measure of causal thinking was obtained from a checklist of thoughts. Ten thoughts (eight causal, two filler) were presented to subjects, with instructions to check off those that had occurred to them while examining the data. An example of a causal thought is, "How good firefighters may take only certain kinds of risks." The measure of causal thinking was the number of causal thoughts checked by the subject. The third measure of causal thinking was the subject's answer to the question, "How much time did you spend thinking about how or why risk preference might cause a person to be successful (or unsuccessful) as a fire fighter?" Subjects indicated their answers on a 9-point rating scale with "1" representing "no time" and "9" representing "very much time." These three measures were followed by the standard debriefing, as in the first study. (3) The final difference in procedure was the measure of subjects' personal beliefs about the true relationship. In the present study, only the measure based on subjects' predictions of success or failure for five new case histories was used. Briefly, five case histories were presented, including responses to risk preference items. The success and failure predictions for each of these cases were scored on the basis of being congruent with a positive relationship (+1) or with a negative one (−1), thus making the range of possible scores +5 to −5.

Results and Discussion

Manipulation Check

Results from the manipulation check showed, as in Experiment 1, that Positive Relationship subjects discovered a positive relationship, $M = 31.96$, $t(46) = 4.87$, $p < .0001$, while Negative Relationship subjects discovered a negative one, $M = -21.72$, $t(46) = 4.87$, $p < .0001$. Also as in Experiment 1, the magnitude of discovered theories was slightly (but nonsignificantly) greater for the abstract than for the concrete conditions, $F(1, 46) = 1.42$, ns.

Dependent Variables

There were no consistent or meaningful effects of subject sex; therefore, all presented means collapse across this variable.

Causal thinking. Three measures of causal thinking were taken to examine the extent to which subjects spontaneously engage in causal processing and to test the prediction that such processing occurs to a greater extent with concrete data. The results from these measures, shown in Table 3, confirm that even when not explicitly asked to explain the relationship data, subjects frequently do so, and do so to a greater extent when examining concrete data than when examining abstract data. For example, 72% of the concrete data subjects listed at least one thought that was of a general causal nature, compared to only 40% of the abstract data subjects. These rates are remarkably similar to (and not significantly different from) those obtained in Experiment 1 from the examination of the written explanations; 81 and 33% of concrete and abstract data subjects wrote general causal explanations when explicitly asked to do so. Thus, the failure to increase theory perseverance in that experiment by instructing

some subjects to write causal explanations can easily be understood as a failure to increase the proportion of subjects who engaged in causal processing ($\chi^2 < 1$ for both the abstract and the concrete comparisons).

As shown in Table 3, concrete data subjects also claimed more causal thoughts on the check list, and spent more time in causal thought than did abstract data subjects, t's(46) = 4.84 and 2.34, p's < .0001 and .03, respectively.

Personal beliefs. Recall that the subjects' personal beliefs about the true relationship were assessed after the debriefing made it clear that the initial data were fictitious. This measure was based on subjects' predictions to five new case histories. Covariance analyses were used, within the positive and negative relationship conditions, to adjust for differences in initial beliefs.

Three basic questions can be addressed by these data. First, was there a significant level of theory perseverance? Second, was the perseverance effect stronger in the concrete than in the abstract conditions? Third, did subjects who spontaneously listed at least one causal thought show more perseverance than those who did not? The relevant ANCOVA yielded affirmative answers to all these questions. First, despite the total discrediting of the initial relationship data, subjects who had initially received positive relationship information continued to believe in a positive relationship, $M = .72$, while negative relationship subjects continued to believe in a negative one, $M = -1.20$. This main perseverance effect was highly significant, $F(1, 33) = 10.24$, $p < .005$. Second, the difference in final beliefs between positive and negative relationship conditions was greater in the concrete than in the abstract data conditions, $F(1, 33) = 3.85$, $p < .07$. That is, the perseverance effect was stronger in the concrete conditions. Finally, subjects who listed at least one general causal thought

TABLE 3
AMOUNT OF CAUSAL PROCESSING AS A FUNCTION OF TYPE OF DATA, EXPERIMENT 2

Measure of causal thinking	Abstract data	Concrete data	Significance test
Thought listing: % subjects who listed at least 1 causal thought	40%	72%	$\chi^2(1) = 5.19$*
Number of causal thoughts checked	2.5	5.0	$t(46) = 4.84$**
Amount of time spent on causal thinking[a]	4.1	5.5	$t(46) = 2.34$*
n	25	25	

[a] Based on a 9-point rating scale, 1 = no time, 9 = very much time.
* $p < .03$.
** $p < .001$.

showed significantly more perseverance than those who listed no causal thoughts, $F(1, 33) = 6.18$, $p < .02$.[4]

As a final test of the causal processing hypothesis, a correlation was calculated between a causal processing index (sum of the z scores on the three variables in Table 3) and degree of theory perseverance. This latter measure was simply the difference between a subject's final belief score and the mean belief score for that subject's relationship condition (positive or negative). The significant correlation that resulted, $r = .34$, $p < .02$, further supported the proposition that increased causal processing increases theory perseverence.

CONCLUSIONS

Several general conclusions may be drawn from these data. First, unwarranted levels of theory perseverance can occur even when the initial theory is based on pallid, abstract data, as demonstrated in Experiment 1. Paradoxically, the perseverance bias is more pervasive when based on the type of data that is most likely to be challenged and discredited—weak but vivid, concrete case history data. Subjects exposed to only two case histories of dubious representativeness clung to their initial theories to a significantly greater extent than did subjects exposed to raw data and statistical summaries of twenty cases. This difference in theory perseverance occurred both immediately after the initial data were discredited (in both experiments) and after a 1-week delay period (Experiment 1).

The importance of the results from the delayed measures in Experiment 1 become clearer when we consider the mechanisms proposed to underlie theory perseverance. Our beliefs about relationships between social variables are probably based on some kind of availability heuristic. We believe that "A" leads to "B" to the extent that we can easily imagine how or why "A" should cause "B", or to the extent that we can easily construct a scenario in which "A" leads to "B". Earlier research had looked at theory perseverance only in a short term situation, however, resulting in the possibility that the manipulations affected theory availability only in short term or working memory. Experiment 1, however, conclusively demonstrated long term (1 week) theory availability effects.

The present studies also showed that causal processing is an important factor in theory perseverance. Furthermore, the more extreme perseverance that occurred in the concrete data conditions seemed to result from the increased propensity to engage in causal processing when examining concrete data. Now if causal processing were unusual, or if important

[4] One reviewer suggested a similar internal analysis be conducted on the perseverance data from Experiment 1, using the categorizations of the written explanations (causal vs not causal) as a factor. Unfortunately, several extremely small cell sizes (i.e., n's = 1 or 2) resulted from that analysis, making it too unreliable to be of any value.

decisions were only rarely based on weak or contradictory data, the implications of these demonstrations of theory perseverance would be less profound. But the present research and other work has shown causal analysis to be a common cognitive process that is likely to be engaged in whenever important, unusual, or surprising events occur (Heider, 1958; Kelley, 1967, 1973; Pyszczynski & Greenberg, 1981; Wong & Weiner, 1981). In addition, many important decisions must be made on the basis of weak or contradictory data that are also full of rich, concrete information (Abelson, 1976; Chapman & Chapman, 1969; Janis & Mann, 1977). The importance of understanding belief perseverance biases is clear at both the individual and the societal level. An important direction for future research is to examine ways of counteracting or reducing belief perseverance. For example, Anderson (1982) has shown that inducing people to create explanations for both possible relationships between a pair of variables reduces the amount of theory perseverance. Further investigation of perseverance biases and of debiasing techniques is clearly warranted.

REFERENCES

Abelson, R. P. Script processing in attitude formation and decision-making. In J. S. Carroll & J. W. Payne (Eds.), *Cognition and social behavior*. Hillsdale, N.J.: Erlbaum, 1976.

Allport, G. W. *The nature of prejudice*. Reading, Mass.: Addison-Wesley, 1954.

Anderson, C. A. Inoculation and counter-explanation: Debiasing techniques in the perseverance of social theories. *Social Cognition*, 1982, **1**, 126–139.

Anderson, C. A., Lepper, Mark 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.

Borgida, E., & Nisbett, R. E. The differential impact of abstract vs concrete information on decisions. *Journal of Applied Social Psychology*, 1977, **7**, 258–271.

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.

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.

D'Agostino, P. R., O'Neill, B. J., & Pavio, A. Memory for pictures and words as a function of level of processing: Depth or dual coding? *Memory and Cognition*, 1977, **5**, 252–256.

Enzle, M. E., Hansen, R. D., & Lowe, C. A. Humanizing the mixed-motive paradigm: Methodological implications from attribution theory. *Simulation and Games*, 1975, **6**, 151–165.

Hamill, R., Wilson, T., & Nisbett, R. E. Insensitivity to sample bias: Generalizing from atypical cases. *Journal of Personality and Social Psychology*, 1980, **39**, 578–589.

Heider, F. *The psychology of interpersonal relations*. New York: Wiley, 1958.

Hovland, C. I., Janis, I. L., & Kelley, H. H. *Communication and persuasion*. New Haven, Conn.: Yale Univ. Press, 1953.

Janis, I. L., & Mann, L. *Decision making: A psychological analysis of conflict, choice, and commitment*. New York: Free Press, 1977.

Kelley, H. H. Attribution theory in social psychology. In D. Levine (Ed.), *Nebraska symposium on motivation* (Vol. 15). Lincoln, Neb.: Univ. of Nebraska Press, 1967.

Kelley, H. H. The processes of causal attribution. *American Psychologist*, 1973, **28**, 107–128.

Luchins, A. S. Experimental attempts to minimize the impact of first impressions. In C. I. Hovland (Ed.), *The order of presentation in persuasion*. New Haven, Conn.: Yale Univ. Press, 1957.

Parker, J. F. The locus of imagery effects in free recall. *American Journal of Psychology*, 1981, **94**, 113–124.

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.

Reyes, R. M., Thompson, W. C., & Bower, G. H. Judgmental biases resulting from differing availabilities of arguments. *Journal of Personality and Social Psychology*, 1980, **39**, 2–12.

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 Univ. Press, 1982.

Ross, L., & Lepper, M. R. The perseverance of beliefs: Empirical and normative considerations. In R. A. Shweder & D. Fiske (Eds.), *New directions for methodology of behavioral sciences: 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 attributional 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 expectation: Effects of real and hypothetical explanations on subjective likelihood. *Journal of Personality and Social Psychology*, 1977, **35**, 817–829.

Smith, C. D. Recognition memory for sentences as a function of concreteness/abstractness and affirmation/negation. *British Journal of Psychology*, 1981, **72**, 125–129.

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

Wong, P. T. P., & Weiner, B. When people ask "why" questions, and the heuristics of attributional search. *Journal of Personality and Social Psychology*, 1981, **40**, 640–663.