ELSEVIER

Cognition 54 (1995) 299-352

COGNITION

The role of covariation versus mechanism information in causal attribution

Woo-kyoung Ahn[a,*], Charles W. Kalish[b], Douglas L. Medin[c], Susan A. Gelman[d]

[a]*Department of Psychology, 317 Life Sciences Building, University of Louisville, 2301 South 3rd Street, Louisville, KY 40292, USA*
[b]*University of Wisconsin at Madison, Madison, WI 53706, USA*
[c]*Northwestern University, 2029 Sheridan Road, Evanston, IL 60208, USA*
[d]*University of Michigan, Ann Arbor, MI 48109, USA*

Received May 15, 1993, final version accepted June 16, 1994

Abstract

Traditional approaches to causal attribution propose that information about covariation of factors is used to identify causes of events. In contrast, we present a series of studies showing that people seek out and prefer information about causal mechanisms rather than information about covariation. Experiments 1, 2 and 3 asked subjects to indicate the kind of information they would need for causal attribution. The subjects tended to seek out information that would provide evidence for or against hypotheses about underlying mechanisms. When asked to provide causes, the subjects' descriptions were also based on causal mechanisms. In Experiment 4, subjects received pieces of conflicting evidence matching in covariation values but differing in whether the evidence included some statement of a mechanism. The influence of evidence was significantly stronger when it included mechanism information. We conclude that people do not treat the task of causal attribution as one of identifying a novel causal relationship between arbitrary factors by relying solely on covariation information. Rather, people attempt to seek out causal mechanisms in developing a causal explanation for a specific event.

1. Introduction

Over the course of experience we are often required to make judgements about the causes of events. Different models of causal inference (or

* Corresponding author.

SSDI 0010-0277(94)00640-7

"attribution") have emphasized different types of information, such as covariation (e.g., Cheng & Novick, 1990a, 1990b, 1991, 1992; Kelley, 1967, 1973), pragmatically constrained information (e.g., Hilton, 1990), temporal contiguity (e.g., Bullock, Gelman, & Baillargeon, 1982; Einhorn & Hogarth, 1986), spatial contiguity (Michotte, 1946/1963), and knowledge about underlying mechanisms (e.g., Bullock et al., 1982; Shultz, 1982; Shultz, Fisher, Pratt, & Rulf, 1986). These studies tend to focus on information use; for example, how information about covariation between a factor and an effect is *used* to identify the factor as the cause. However, just because people can use a particular kind of information does not necessarily mean that they *typically* do so. People may be able to carry out analyses of covariation, but that may not be the process they spontaneously use. In order to judge the methods of causal attribution people use, we examined what kinds of information subjects request when asked to identify the causes of some events. We believe that information search strategies may shed light on causal inference processes.

Very few studies have directly examined information search associated with causal attribution. The majority of studies have focused on the *use* of a particular kind of information. In the next section of this article, we describe covariation models of causal attribution and extend them to predictions about information search. Then we outline a complementary approach emphasizing the role of knowledge for underlying mechanisms. The debate between the two approaches has a long history in philosophy (See Kitcher & Salmon, 1989, for a review); the statistical relevance view says that correlational facts constitute causal connections (e.g., Salmon, 1970; Suppes, 1970) whereas the deterministic view or the power view of causation claims that causes have to determine their effects by necessity (e.g., Harré & Madden, 1975; Pappineau, 1989; Salmon, 1984). The debate has more recently received attention in psychology (e.g., Cheng, 1993; White, 1990). These two approaches to points of view will provide the context for our studies of information-seeking.

2. Covariational approach

Research on causal reasoning has concentrated on the question of how people decide *who* or *what* caused an effect to occur (as opposed to how the event occurred). Given an event such as John having a car accident on Route 7 yesterday, the question for traditional attribution theory is how people decide, for example, whether something about John, something about the road, or something about yesterday caused the accident. Here the issue is "who or what caused the event?" In other words, which of the components (John, Route 7, yesterday) that make up the event should be identified as the cause?

In answering this type of "who/what" question, a large number of models (e.g., Cheng & Novick, 1990a, 1990b, 1991, 1992; Hewstone & Jaspars, 1987; Kelley, 1967, 1973) have focused on one particular form of reasoning: the covariation principle. According to this principle, the process of identifying a cause involves comparing an event with other, contrasting events. The pedigree of this approach can be traced back to Mill's method of differences (Mill, 1843/1973). The general idea is that the pattern of elements that covary with the effect will identify a cause. That is, "the effect is attributed to that condition which is present when the effect is present and which is absent when the effect is absent" (Kelley, 1967, p. 194).

A concrete example will help clarify the process. Consider again John's car accident. If one want to find out what caused this event, the covariation approach will start out with possible candidate factors responsible for the accident, such as John and Route 7. Then the method of differences would proceed by looking at the covariation between these factors and other events. For instance, if other people also had accidents at that time, then the event might not be due to something special about John. If John tends not to have car accidents on other occasions, then the event might be due to something special about the particular situation. By sifting through the combinations of factors that go along with car accidents, one can assign responsibility to some set of the factors in the target event. The following two sections describe two examples of the covariational approach, which will then be contrasted with an alternative approach in our experiments.

2.1. Kelley's model

Kelley proposed a psychological model of causal attribution based on the principle of covariation. His model is characterized as an ANOVA model because it involves a process analogous to an analysis of variance. In this model, an event is represented in terms of person, stimulus, and occasion dimensions. Causal attributions are made by contrasting the target event with other events along each of the three dimensions. In the event, "John had a traffic accident last night", the important information variables are consensus on the person dimension ("Did other people also have a traffic accident last night?"), distinctiveness on the stimulus dimension ("Did John have other kinds of accidents last night?"), and consistency on the occasion dimension ("Does John usually have traffic accidents?"). According to the model, the configuration of these three information variables determines whether person, stimulus, or occasion attribution will be made (Kelley, 1967, 1973).

The model has been supported by a great deal of experimental data (e.g., McArthur, 1972; see Kassin, 1979, for a review) but these experiments also showed some biases not explained by the model. For example, McArthur (1972) and many others (e.g., Jaspars, Hewston, & Fincham, 1983; see Cheng & Novick, 1990a and Hewstone, 1989, for a review) have found that people show a bias favoring person attributions over stimulus and occasion attributions.

2.2. Probabilistic contrast model

In order to explain this apparent bias, Cheng and Novick developed a probabilistic contrast model (1990a, 1990b, 1991, 1992). Whereas Kelley's model focuses on only three particular dimensions (person, stimulus, and occasion), the probabilistic contrast model is not restricted to particular sets of dimensions. It also considers interaction effects among the dimensions, whereas Kelley's consistency, consensus, and distinctiveness concern only the main effects. We will briefly describe this model and use it as our paradigmatic exemplar of the covariation approach.

The model assumes that people start out with some pre-existing conceptions of possible causal factors (e.g., John, John's car, Route 7 in the above example). These factors might be "psychologically prior" factors determined by innate biases or previous knowledge of underlying causal mechanisms (Cheng & Novick, 1992). Then contrasts (or covariation) are computed for "focal sets" that are restricted to events in which each of these factors is present and absent. If the proportion of events for which the effect occurs given presence of the factor is noticeably different from the proportion of events for which the effect occurs given absence of the factor, the factor is inferred as the cause of the event.

For example, consider an event taken from Cheng and Novick (1990a), (1) "Beth said the morning prayer on this occasion", and the following seven pieces of covariation information when the factors under consideration are person, stimulus, and occasion. (2) Everyone else said the morning prayer on this occasion. (3) Beth did not say any other prayer on this occasion. (4) On all other occasions, Beth always had said the morning prayer. (5) On all other occasions, everyone else said the morning prayer. (6) On all other occasions, Beth said all other prayers. (7) Nobody else said all other prayers on this occasion. (8) On all other occasions, everyone else said all the prayers. This situation can be represented as in Fig. 1 where the target event corresponds to cell 1 and each of the above pieces of covariation information corresponds to the cell indicated by the same number. According to the probabilistic contrast model, the normative causal attribution given this combination of covariation information is to choose that there is something special about the combination of the morning prayer and the occasion (see Cheng & Novick, 1990a, for the proof).

Note that Kelley's model needs only cells 1, 2, 3, and 4 for covariational analysis whereas the probabilistic contrast model requires all eight cells assuming that the person, the occasion, and the stimulus are the dimensions of interest. The probabilistic contrast model argues that the apparent bias toward person attributions observed in other experiments (e.g., McArthur, 1972) is due to the subjects' implicit assumptions about covariation not

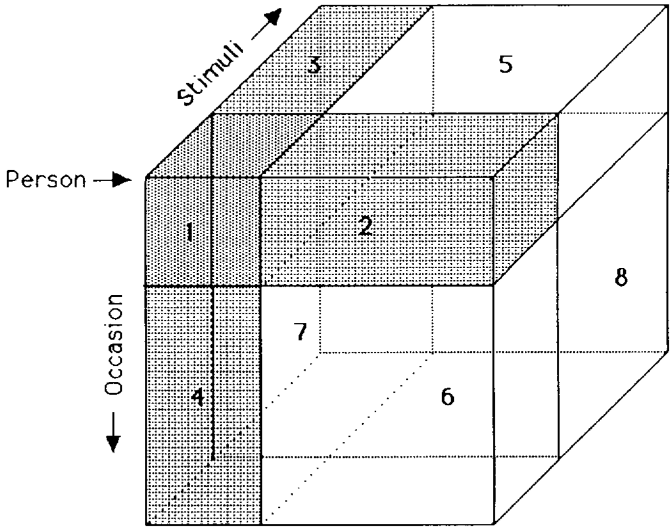

Fig. 1. The eight regions of information for Person × Stimuli × Occasion in the probabilistic contrast model.

specified by the model (i.e., cells 5, 6, 7, and 8). In Cheng and Novick's (1990a) experiments, the subjects' task was to identify the cause of the target event given the full set of covariation information.[1] They used a forced-choice task where the subjects could attribute the event to person, stimulus, occasion, the three two-way interaction effects, and/or the one three-way interaction effect. When the full information was given, the subjects responded in a normative manner as predicted by the probabilistic contrast model.

3. Mechanism approach

A complementary approach to causal inference stresses the role of mechanisms or processes underlying an event: the crucial determinant of inferring causality is people's knowledge of or belief in the existence of mechanisms.[2] A mechanism is some component of an event which is thought to have causal force or causal necessity. Rather than seeing events as composed of a set of independent factors (person, stimulus, occasion), this approach treats events as composed of "surface" factors (John, Route 7)

[1] In the actual experimental materials of Cheng and Novick (1990a), some of the sentences from (2) to (8) were combined.

[2] The current study does not deal with the metaphysical distinction between mere covariation and causal necessity. Instead, it deals with psychological phenomenon of inferring causality. That is, why people believe a relationship to be causal, not how to tell whether the relationship actually is or is not. (See later sections for more detail.)

and one (or more) underlying, responsible mechanisms (e.g., John's drunkenness, the failure of the brakes). Factors (such as John) may be identified as the cause of an event via the discovery of some mechanism rather than covariational analyses.

When seeking the cause of John's accident, the focus of the mechanism analysis would be on discovering the *process* underlying the relationship between the cause and the effect. For instance, people might ask such questions as "Was John drunk?" or "Was there a mechanical problem with the car?" Here the search for a cause proceeds by identifying the mechanism or means that were operating during the event. We can describe this approach as a focus on "how" questions; how did the event happen? Unless this question is answered, people would be reluctant to accept a factor as the cause of the event.

In contrast, covariation models do not answer questions about how the *factor* caused the event but simply assign a particular factor (or combination of factors) as the cause of the event (Hewstone, 1989). Consider our example of John's car accident. Suppose, after looking at the relevant comparison events, we discover that John is particularly susceptible to car crashes (e.g., John has had many car crashes but other people have had very few). We might well assume that something about John was the cause of his most recent crash. But what about John? Is he an alcoholic, a professional race track driver, a bad driver, or is there someone out to get John who keeps arranging these accidents? Although covariation information allows us to identify a factor, we do not know the nature of the connection between the factor and the effect. One question the current experiments attempt to answer is what type of information (process or factors) people seek out during causal inference.

4. Comparing covariation and mechanism approaches

In this section, we will describe three differences between the covariation and the mechanism approaches: (a) *types of vocabularies* which leads us to predict different levels of analysis; (b) causal attribution as inferring *a necessary connection* versus inducing *a regularity* which leads us to focus on preconditions in the same event rather than interest in other combinations of the same factors; and (c) *content specificity* of attribution which leads us to predict that process will differ depending on content of the events.

4.1. *Types of vocabularies*

The two approaches differ in the types of vocabularies used in their explanations. The mechanism-based explanation characterizes the event at a different level of abstraction and involves the introduction of a new set of entities or processes. Suppose John had a car accident and the underlying mechanism was engine malfunction. Here we have introduced new pieces of vocabulary (engine, malfunction) that are not simple re-descriptions of the elements already described in the event. Gopnik and Wellman (1994) describe this as a distinction between "empirical generalizations" and "theoretical entities". Empirical generalizations are "orderings, partitionings, and glosses of evidence and experience but ones that share the same basic vocabulary as the evidence itself". For example, one may categorize plants into ones with no stems, ones with green stems, and ones with woody stems based on the surface features of the plants. On the other hand, theories include theoretical constructs to provide causal, explanatory level of analysis of evidence. Therefore, the vocabularies involved in theories refer to entities removed from and underlying the evidential phenomena themselves.

In our framework, empirical generalizations correspond to the covariation approach whereas theoretical entities correspond to the mechanism approach. The provision of a mechanism turns an empirical generalization, a statement of a regularity, into an explanation, an account of that regularity. In contrast, the causes identified by the covariation principle are at the same level of abstraction as the event described. As Cheng and Novick stated (1990a): "Contrasts are assumed to be computed for attended dimensions that are present in the event to be explained" (p. 549). That is, if the event was described as involving John, Route 7, a day (yesterday), and an accident, it is the relationship among these terms (at this level of generality) that people are trying to discover. The result of the attribution will also be phrased at this level; for instance, as an assertion of a positive relationship between John and car accidents (e.g., "something special about John caused the car accident").

Of course, the event descriptions can include more specific, mechanism-based constructs and one might argue that covariation analysis can eventually generate mechanism-based explanations when causal candidates are mechanisms. For example, given the event "John, who was *drunk*, had a traffic accident last night", one might attribute drunk driving as the cause of the accident through some covariation analysis. This explanation seems inappropriate when the event is described at that level (e.g., "The drunk had an accident because he/she was drunk").[3] When asked what caused the *drunk* to have a car accident, one seems to be asking for the explanation beyond the level of the event description (e.g., he/she was disoriented, his/her reaction time was too slow to make the curve). The emphasis of the mechanism approach is on this tendency to go beyond the surface level of the event description. Whether or not people do so is the empirical question to be tested in our experiments.

[3] In our discussion of the mechanism approach, we treat causal attribution as being like explanations.

4.2. Causal necessity versus inductive justification

According to the mechanism approach, the causal relationship in people's mind is not one of inductive generalizations based on probabilistic contrast but rather one of causal necessity (see Harré & Madden, 1975; Salmon, 1984, for philosophical discussion on causal necessity). The mechanism-based explanations entail causal necessity or determinism because they involve theoretical constructs or principles that universally hold as long as their preconditions are met. Therefore, the mechanism approach would argue that when people attempt causal attributions, they may seek out information about whether the target event satisfies the preconditions for a hypothesized mechanism. Once the preconditions are met, one can be certain that the mechanism was at work and it was the cause of the event because of the necessary relationship.

On the other hand, a covariation analysis does not truly demonstrate that a particular factor is necessary or sufficient for an effect, only that there has been a reliable relationship in the past (i.e., inductive justification). Therefore, what is crucial for the covariation strategy is to increase the probability of this belief by seeking out more evidence for the covariation: how the factors in the particular target event covary with other events.

When people are trying to decide, for example, whether Kim is the cause of the traffic accident or not, according to the mechanism approach they would ask for information about conditions in which the particular accident occurred rather than information about whether Kim is reliably associated with traffic accidents in general. The latter strategy only gives us inductive justification to increase one's probabilistic belief that Kim caused the particular accident. The former provides necessity. The covariation information might be useful in suggesting candidate mechanisms (see General discussion) but it does not give a satisfying answer for specific mechanisms involved in the particular target event. Experiments 1, 2, and 3 test what type of information people spontaneously seek out with respect to this dimension.

Due to the lack of the necessary relationship, the covariation approach has difficulty distinguishing between regular covariation and causality (Hewstone, 1989; Hilton, 1988; Shaver, 1981; but see Spirtes, Glymour, & Scheines, 1993, for more recent attempts). A covaration between two factors can be simply due to a spurious correlation manifested by a third factor. For example, when an increase in the number of churches covaries with an increase in the number of prostitutes, it is not necessarily because the former causes the latter (or vice versa) but maybe because both factors are caused by an increase in population. We believe that this problem arises because the causality according to the covariation approach is not based on the causal necessity.

The mechanism approach, on the other hand, can easily handle this problem by treating any covariation lacking mechanisms as potentially spurious or coincidental. For example, people would have difficulty believing that an increase in churches causes an increase in prostitution not because of lack of confidence in covaration but because of lack of knowledge about underlying mechanisms. This leads us to predict that people will search for preconditions or more information about the target event.

4.3. Content-specific process versus general strategy

Previously, the mechanism approach has emphasized the role of the background knowledge.[4] We assume that generally people know a set of mechanisms and during the process of causal attribution they try to discover whether one was operating in the particular situation. Previous knowledge about mechanisms can be directly applied to a specific event or adapted to an event through analogical reasoning. In the previous example concerning John's traffic accident, the knowledge about general mechanisms (e.g., how drunkeness affects motor skills) helps people determine the causal status of relevant factors (e.g., John). Or the knowledge that some county was unable to maintain its parks might lead one to suspect that Route 7 was also not maintained in an acceptable condition. Either way, hypotheses are generated from previous knowledge, and the way they are tested will depend upon the particular content of the event to be accounted for. Discovering the cause of a car accident may involve different hypotheses and methods than searching for the cause of an airplane accident. The possible causes of John's accident may be different than Julie's. In other words, the causal attribution process is content-specific.

In contrast, the covariation approach suggests that people take the same, general approach to all causal attribution. This approach has previously emphasized causal induction; the task of causal inference is to induce a new causal relationship (e.g., finding out that there is covariation between John and traffic accidents). The particulars of the event are treated as uninterpreted variables. The specifics of an event may affect which factors are considered as possible causes (e.g., ruling out hair color as a potential cause of accidents). The role of particular knowledge about content in the attribution process is often seen as introducing "bias" (Cheng & Novick, 1990a). Analysis of covariation is essentially a syntactic process tabulating the occurrence or non-occurrence of a factor. Thus determining the cause of "X had a Y at Z" would not differ from determining the cause of "John had

[4] Another knowledge-based approach to causal attribution has been presented by Abelson, Lalljee, and their colleagues (Abelson & Lalljee, 1988; Lalljee & Abelson, 1983; Leddo & Abelson, 1986), which is often called "knowledge-structure approach". This approach is different from the mechanism approach because the former is a content-oriented approach where particular contents of existing schemata (e.g., visit-to-doctor's office schema) are used for causal attribution (Hewstone, 1989). On the other hand, the mechanism approach is based on more general schemata or concepts. In addition, the types of schemata on which the knowledge-structure approach has focused do not necessarily involve mechanisms.

a car accident last night". The covariation approach suggests some very modest initial hypotheses (e.g., "It was something about John") which are then tested through covariation analysis.

The covariation approach so far is more appropriate for the cases where one acquires new information about how two factors covary. This approach has been weak about how prestored covariation information would apply to novel situations. For us, people do not very often start from scratch and go looking for novel causal relations. Instead they prefer to look at new situations as instances of familiar causes. We are not debating whether or not covariation reasoning ever happens, but we are arguing that it does not happen very often when people are looking for the cause of an event.

4.4. Other related issues

We have discussed that causal attribution process involves identifying mechanisms that are not specified in the event descriptions. One might wonder then at which level the pursuit for mechanism-based explanation stops. A notorious problem with the mechanism-based explanations is that when pressed for further elaboration of a mechanism, people eventually fall back on covariation for its grounding. If one asks why enough, eventually one gets to "It just does. Whenever X happens, Y happens" (see Spirtes et al., 1992 – an attempt to explain mechanism information in terms of higher-order correlation). Mechanism does, in fact, entail covariation in the sense that "causal force" is associated with the belief that in the absence of the mechanism the outcome would not have occurred (unless an alternative cause were present).

Although the current study does not focus on this issue, there are at least two ways of dealing with this problem. First, presumably conversational maxims and pragmatics may determine the exact level at which a mechanism is stated (e.g., Hilton, 1990). In addition, people's background knowledge about abnormal conditions can also determine the level appropriate for the context (e.g., Hart & Honoré, 1959; Hilton & Slugoski, 1986; Mackie, 1975). Second, our speculation is that the irreducible or bottomed-out covariation implies "there must be a reason but I don't know it" rather than belief in mere probabilistic covariations or associations. (See Malt, 1990, for similar findings in categorization.) Several cognitive psychologists (Gelman & Wellman, 1991; Medin & Ortony, 1988) have made a similar proposal regarding concepts. People's representations of things might reflect a belief that things have essences although they do not know what the essences are. They simply believe that there is some causal basis for perceived regularities even when they do not know the causal basis. When presented with such "bottomed-out" correlational findings, people may fill in with a "mechanism place-holder", an expectation that there is some more fundamental level underlying the relationship of cause and effect.

A related criticism is that mechanisms cannot be pitted against covariation because mechanism information eventually is built upon covariation. One might wonder how the mechanism information is acquired. For example, in Kim's car accident we might identify engine malfunction as the mechanism causing the accident. Then the question is how we learn that engine malfunction causes car accidents. Didn't they have to come from covariation? This question of acquisition can only be answered through developmental studies. (But see the following section for the support from developmental studies for the mechanism approach.) Our study rather focuses on whether people use these beliefs about mechanisms wherever they come from.

In addition, one might wonder what constitutes the mechanism information. Isn't it just some collection of covariation information? This question seems also inappropriate for the current study because we do not try to give a metaphysical account of mechanism as Hume did (1978); we attempt to examine a psychological phenomena: how people perceive and use mechanisms in causal inference rather than what causality or mechanism actually is. In short, we believe that covariations without the support of mechanisms are *psychologically* different from mechanisms and that people prefer explanations involving mechanism that go beyond the event description.

4.5. Summary

We have outlined two contrasting approaches to the identification of the causes of events. The mechanism approach focuses on the generation and testing of hypotheses about underlying mechanism-based causes. Based on the discussion presented so far, we can define a mechanism as follows; The mechanism information in causal reasoning specifies through which *processes* the event *must* have occurred (i.e., how a factor led to the consequence) by using vocabularies describing entities that are *not* presented in the event descriptions. Alternative models stress a general process involving the analysis of covariation between factors and effects. (See Table 1 for a summary of the distinction.) The question then is which method (covariation or mechanism analyses) *typifies* people's approach to causal attribution? In the following section, we briefly review some evidence for the two approaches and then suggest our own test based on the process of information search.

5. Previous studies on the importance of mechanisms

A number of studies have suggested the importance of the presence or absence of a mechanism in judging a causal relationship. Subjects' causal attributions are influenced by whether a posited relationship between two factors can be seen as an instance of a known causal relationship or whether

Table 1
Summary of distinction between the mechanism approach and the covariation approach

Dimensions	Mechanism approach	Covariation approach
Types of vocabularies	A new set of entities or processes that are not specified in the event descriptions	Same level of vocabularies as the event descriptions
Causal necessity vs. induction	Causality based on causal necessity	Causality based on inductive, probabilistic generalizations
Content effect	Content-specific process	Content general process
Correlation vs. causality	Fundamentally different	Cannot discern

support for the posited relationship comes solely from covariation information. Koslowski, Spilton, and Snipper (1981) have shown that covariation of a factor with an effect results in only a weak causal attribution when subjects are not presented with a mechanism for the relationship. Conversely, the presence of a mechanism in the absence of covariation leads to a moderate causal attribution. Shultz (1982, 1986) has argued that children will identify a cause with a factor that has a mechanism rather than a factor that merely occurs closer in space and/or time to the effect (see also Bullock et al., 1982).

Prior knowledge on mechanisms was found to affect judgement of correlation. Chapman and Chapman (1967, 1969) presented evidence that therapists perceived correlations between test results and psychological disorders when, in fact, there were none. These results were interpreted as being due to people's prior knowledge on how the cause (i.e., disorders) might lead to the effect (i.e., test results). In concept learning, Pazzani (1991) also showed that prior knowledge about causal mechanisms reversed the difficulty of learning conjunctive versus disjunctive concepts. Thus, it seems clear that information about underlying mechanisms is an important component of causal reasoning and concept learning.

6. Information-seeking in causal attribution: overview of experiments

One natural source of evidence regarding the way people identify the causes of events would seem to be the information that they spontaneously seek out in making their judgements. It seems to be a natural task because in many real world situations, initial information is sketchy and people have to explore a variety of avenues to develop an explanation. Previously, only a few studies have used an information-search paradigm. Furthermore, these studies do not say much about various types of explanations (i.e., covariation-based or mechanism-based) because they have focused on the question of which kinds of covariation information (i.e., consensus, consistency, or distinctiveness) subjects look for (e.g., Major, 1980). In contrast, Lalljee, Lamb, Furnham, and Jaspars (1984) examined various types of questions subjects asked, such as specific hypothesis-testing questions, and found very few requests for information about covariation. Lipe (1985) focused on the use of counterfactual information (i.e., whether the event would have occurred if the proposed cause had not occurred) but has found that only one out of 89 generated questions was counterfactual and that most of the responses were questions about specific alternative explanations. These experiments used event descriptions that are quite different from the ones used in experiments supporting the covariation approach and subjects seem to have rich background knowledge about these event descriptions.

The present experiments examine the kinds of information people spontaneously seek out when asked to identify causes in various types of events (i.e., both knowledge-rich and poor). The two approaches described so far can serve as frameworks for the types of information that people look for in those situations. However, it should be noted that theoretical approaches to attribution have not directly made predictions about information search. Covariation models so far have dealt with the use of covariation information when that information is available; causal attribution processes might be orthogonal to information-seeking strategies for causal attribution.

6.1. Covariation

As a heuristic for thinking about information search we offer some simplified models representing each approach. A covariation model would run something like this: an explainer starts out with possible causal candidates, which come from an event description. Since covariation information is the most central and necessary information in causal attribution, the first thing that the explainer will seek out is information about the covariation between the candidate factors and the effect (i.e., absence/presence of these factors given absence/presence of the effect).

6.2. Mechanism

A mechanism model leads to different expectations. In order to identify mechanisms, the most useful information would be specific information about the processes underlying the target event rather than information about the way the event covaries with other events. For instance, in determining why John's car came to crash, we would want more specific facts about how the particular crash occurred. In most cases, an explainer would attempt to directly apply his/her knowledge of possible mechanisms to the target event. The explainer will develop hypotheses based on prior knowledge (e.g., John was drunk) and keep on testing them by asking whether the event fits with the hypotheses. So, for example, data might be sought regarding preconditions of the mechanisms, such as John's alcohol consumption that night, the way the car crashed, and so on. These questions on the preconditions that the explainer would ask are about the target event itself and not about other events (i.e., covariation questions). On finding that the preconditions for the hypothesized mechanism are met (e.g., John had a lot to drink), the corresponding hypothesis would be selected as correct.

To sum up: the covariation approach relies on direct information about covariation between elements of an event. The mechanism approach relies on information about how the event was caused. These two approaches predict differences in kinds of information people would ask for during causal attribution. The covariation approach should predict requests for covariation information between factors mentioned in the event. The mechanism approach should predict a search for features of the event that would allow people to test candidate mechanisms.

EXPERIMENT 1

The goal of Experiment 1 was to examine the kinds of information people ask for when searching for the cause of an event. A variety of events was use because differences in background knowledge might lead to different strategies. One subset of events includes those used in a previous study of a covariation model of causal reasoning (Cheng & Novick, 1990a). We used these materials as an interesting ground in light of the fact that they showed that people could use covariation. In addition, these events were selected by Cheng and Novick (1990a) to minimize background knowledge and prior beliefs. For this reason, these sentences will be called "no-schema sentences". In addition, we developed 16 more events from various domains, which can be categorized into good, bad, obvious, and nonsense events. Subjects presumably have rich background knowledge on some of the events (e.g. "The student performed perfectly on the final exam in sociology") but none on such sentences as nonsense ones (e.g. "XB12 mimbled the wug at fulmer"). These sentences will be called "varied-schema sentences".

Two sets of instructions were used. One kind of instruction was especially designed to elicit covariation questions so that we would not miss any covariation-based reasoning if it happened. The other instructions were neutral with respect to the types of causal reasoning.

7. Method

7.1. Design and procedure

Subjects were presented with booklets containing single sentence descriptions of events. Their task was to write down any questions they would want to have answered in order to identify the causes of the events. After finishing all the items, they were also asked to go back to the previous items and to "write down the most likely cause" for each event.

Because experimental instructions could influence the type of causal reasoning strategy, two types of instructions were employed. One type was designed to encourage covariational reasoning and the other was designed to be neutral. Two separate groups of subjects were run. Both groups were told that they would see several statements describing certain events and that their task was to generate questions that they would like the answers to, if they were to identify the causes of the events. They were also told that they might ask as many questions as they liked. The covariation instruction group received additional instructions:

> Each of the events can be thought of as composed of multiple factors. The questions you should ask should try to get at which factor (or factors) was responsible for the event. For example, given the event, "John did not drink wine while camping", you could think of various components that went into causing this event. Maybe it was something about John that led to his not drinking the wine. On the other hand, maybe there was something special about the wine that caused him not to drink it. It could also be that the camping caused him not to drink the wine at the time. Of course, you could also ask questions designed to discover if some combination of factors (e.g., combination of John and wine, combination of wine and camping, combination of John and camping, or combination of John, wine, and camping) caused the event.
>
> In this task, you do not need to ask the questions you would need to discover the specific cause of the event, just which factor is responsible. Thus identifying the wine as the cause is sufficient and you don't need to ask questions to figure out if the wine was bad, or any other specific thing about the wine. Your questions should just attempt to identify the general factor(s) responsible.

7.2. Materials

Two groups received different sets of materials. One group received the 15 no-schema sentences taken from Cheng and Novick (1990a, see Appendix A) and the other group received the 16 varied-schema sentences (see Appendix B). These 16 statements can be grouped into; four statements about good events, four about bad events, four about obvious events, and four nonsense sentences.

7.3. Subjects

There were 27 subjects who were undergraduate students at the University of Michigan, participating in partial fulfillment of course requirements for introductory psychology. Subjects were randomly assigned to one of the four conditions. There were seven subjects who received the no-schema sentences with the neutral instructions, eight who received the varied-schema sentences with the covariation instructions, six who received the no-schema sentences with neutral instructions, and six who received the varied-schema sentences with the covariation instructions.

7.4. Scoring

The questions generated by subjects were first separated into ANOVA versus non-ANOVA questions with respect to the following criteria. If a question passed *any* of the following criteria, it was considered as an ANOVA question.

(a) Does the question ask about other events? For example, for the event "Dave would not eat rabbit meat on this occasion", the question "Does Dave usually ear rabbits?" is an ANOVA question because this question is about what Dave does on other occasions. On the other hand, the question "Is Dave a vegetarian?" is a non-ANOVA question.

(b) Does the question identify only a responsible factor stated in the event? For example, for the event, "The student performed perfectly on the final exam in sociology", the question, "Did other students do well in the final exam?" was classified as an ANOVA question because it simply attempted to identify whether there was something special about the student.

(c) When the question refers to a factor not stated in the event, is it followed by the complete set of other questions necessary for the covariation analysis? Consider the event "Dave would not eat rabbit meat on this occasion". Subjects might select factors at a more specific level, such as the fact that Dave had a pet rabbit. Then one might ask, "Did Dave have a pet rabbit?" This would seem to be a question about a precondition of some mechanism (some sort of psychological process of identification or disgust) but it could just be an element of a covariation analysis. Maybe the subject is proposing a covariation analysis of the factor "has a pet rabbit" with "won't eat rabbit meat". If this were the case, then, they should go on to ask how the factor covaries with the effect (e.g., "Did other people with pet rabbits eat the rabbit meat?" or "Did Dave eat the meat of other animals that he didn't keep as pets?"). If people did not ask these follow-up questions, we assumed that they already knew that keeping a rabbit as a pet was a cause of not eating rabbit meat and were not performing a covariation analysis of this factor. In fact, no one in the current experiment ever followed up in this way.[5] We report this criterion here because this last criterion eliminated any possibilities of excluding ANOVA questions solely on the basis of levels of questions.

Two judges independently scored half of the data. The two judges agreed on 90% of the items in classifying questions as ANOVA questions. To avoid biasing the results toward the mechanism hypothesis, all conflicts were resolved in favor of an ANOVA classification. For example, a question such as "Did Dave have a bad experience with a raccoon?" was treated as an ANOVA question even though, presumably, its intent was to discover a fact about Dave's mental state at the time rather than to elicit information about the frequency with which Dave was afraid of raccoons.

The non-ANOVA questions were further classified into the following three categories: (a) hypothesis-testing questions, the questions asking more about the target event by referring to the mechanisms that are not stated in the event about *how* the event happened (e.g., "Did he have a toothache?" for the event "Dave would not eat rabbit meat on this occasion"); (b) information-seeking questions, the questions asking for further factual information about the target events (e.g., "What is this occasion?"); (c) effect questions, the questions asking about the effect or the consequence of the event (e.g., "How many people died?" for the event, "Flight 921 crashed in Lincoln Nebraska last winter"); and (d) other types of questions. The other types of questions include miscellaneous questions such as questions about linguistic meanings of words (linguisitic questions) (e.g., "What are toves?"), and irrelevant questions (e.g., "Are we supposed to answer this?").

The distinction between the hypothesis-testing questions and the information-seeking questions is not entirely clear because the subjects were responding in natural language and there are all types of conventions about what one can assume and what one need to mention. Consider the questions such as "Was anyone with Jane?" for the event, "Jane had fun washing dishes on this occasion". This question does not contain clear statements of mechanisms but it seems to represent more than bottom-up, information accumulation. For this reason, the raters agreed on the distinction between hypothesis-testing and information-seeking questions only 68% of the time. This finer distinction, however, is independent of our main interest of whether people spontaneously seek out covariation-based information (i.e., ANOVA vs. non-ANOVA distinction). The distinction was made simply to obtain richer data analyses.

[5] At some previous time, the subject might have done a covariation analysis of pet keeping to meat eating. But the important point is that they are not doing the covariation analysis when they make causal attributions.

The explanations generated by subjects were grouped into either ANOVA explanations or mechanism explanations. Mechanism explanations provided a statement of a mechanism which mediated between a causal factor and the effect. ANOVA explanations just stated the causal relationship without providing a mechanism. An example of a mechanism explanation is "John's raccoon is in a cage" for the event, "Sam was not afraid of John's raccoon on this occasion". An ANOVA explanation would be "Sam is never afraid of John's raccoon". Although this ANOVA explanation merely rules out a possible causal factor (the occasion) without identifying the person (Sam) or the stimulus (the raccoon) as the cause, we used a liberal criterion and treated this type of explanation as an ANOVA explanation.

8. Results

8.1. *Effect of types of instructions*

The covariation instructions did not produce more ANOVA questions or ANOVA explanations than the normal instructions: There was no reliable interaction effect between type of instruction and type of question, $p > .50$, and no interaction effect between type of explanations and type of instruction, $p > .50$. Therefore, in all of the following analyses the data from the two instructional groups were combined.

8.2. *Main analysis of questions*

Fig. 2 shows the average of percentages of each question type. Overall, we observed very few ANOVA questions (13.23% of all the questions in response to the no-schema sentences and 6.87% in response to the varied-schema sentences). Most the non-ANOVA questions were coded as hypothesis-testing questions (63.09% in response to the varied-schema sentences and 66.25% in response to the no-schema sentences). For each of the 31 items, there were more hypothesis-testing questions than ANOVA questions. In addition, all subjects but one produced more hypothesis-testing questions than ANOVA questions. (The exceptional subject's questions were mostly information-seeking questions.)

Since the number of items differs between the no- and the varied-schema sentences, the average frequencies of each question type per item were used to test the effect of events (no-schema vs. varied-schema sentences) and type of question. There was no reliable main effect of events, $p > .10$, but there was a reliable main effect of type of question ($F(4, 100) = 58.768$, $p < .01$, $MSE = 0.429$). The interaction between events and type of question was not reliable, $p > .10$.

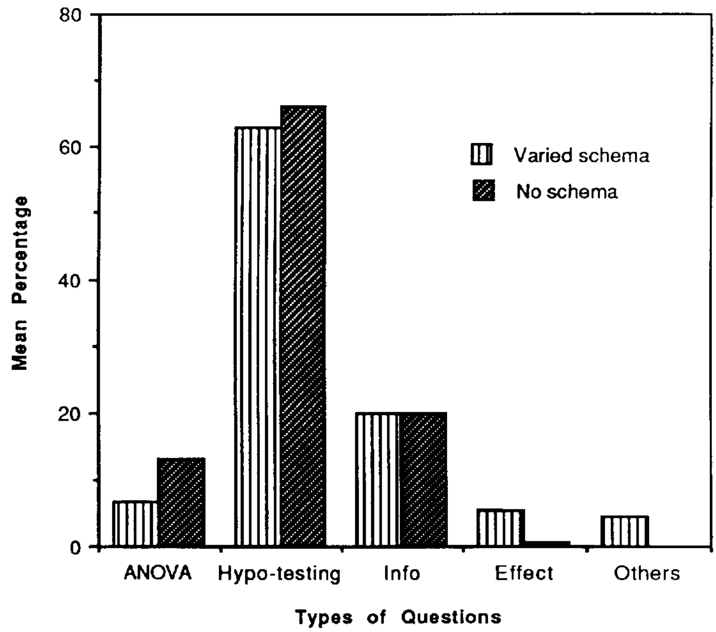

Fig. 2. Percentage of each question type for the varied-schema condition and the no-schema condition in Experiment 1.

8.3. Secondary analyses of questions for each type of sentences

More ANOVA questions were asked from the no-schema sentences (0.51 question per item) than from varied-schema sentences (0.24 question per item), $F(1, 25) = 6.138$, $\text{MSE} = 0.04$, $p < .05$. Of the ANOVA questions, 41% were questions about occasions. The subjects who received the no-schema sentences produced more ANOVA questions about occasions than the subjects who received varied-schema sentences (5.4% for the no-schema sentences and 2.6% for varied-schema sentences). This result might be because the no-schema sentences use the wording "on this occasion" whereas varied-schema sentences specified an occasion. This repetition of "on this occasion" in the no-schema sentences may have called attention to the occasion aspect of events. Participants frequently asked "What was the occasion?"

Further analyses were carried out with varied-schema sentences broken down into four item types. Fig. 3 shows the frequencies of each type of questions in each item type. ANOVA questions were the least frequent regardless of the item type. There were more hypothesis-testing questions than information-seeking questions for each of the items except for the nonsense items.

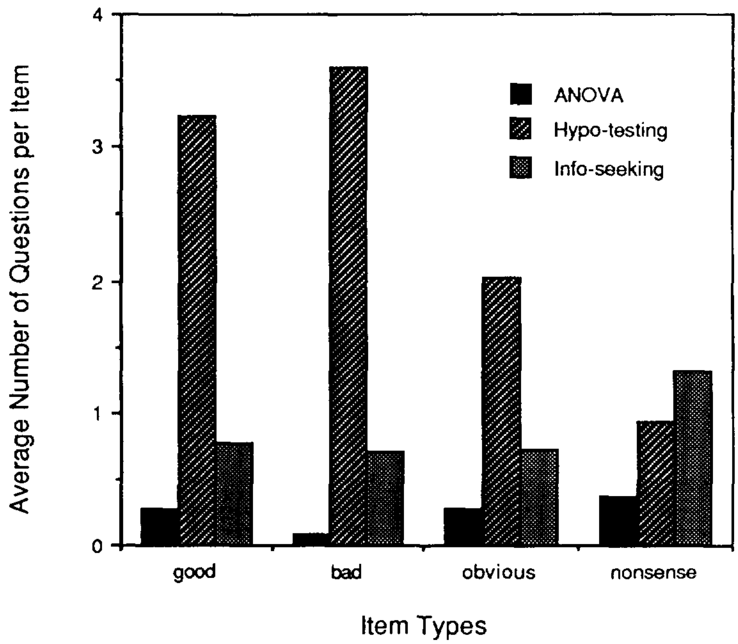

Fig. 3. Average number of questions for different item types for the no-schema condition in Experiment 1.

8.4. Analysis of explanations

After generating questions, the subjects gave the most plausible explanation for each item. Analysis of the explanations showed the same pattern of results as for the questions. Despite our lenient scoring criterion for an ANOVA explanation, the vast majority of explanations (83%)[6] were mechanism rather than ANOVA explanations. In response to no-schema sentences 18% of the explanations were of ANOVA type, and in response to varied-schema sentences 15% were of ANOVA type.

9. Discussion

The results demonstrate that subjects focus on underlying causal relationships in both information requests and explanations. Most questions sought more information about the target event, rather than information about how the stated factors covaried. Subjects most frequently asked questions that

[6] Because only one explanation was given by each subject for each item, the measures of frequency information for all of the following relative measures for the analysis of explanations in Experiments 1, 2, and 3 can be obtained by multiplying each percentage by the number of items.

introduced some new element into the event. For example, presented with the event "John was not afraid of Dave's raccoon on this occasion", many subjects asked whether the raccoon was in a cage or not. Here the subjects seem to attempt to account for the event in terms of a mechanism already known to them (e.g., caged animals cannot attack people and hence there is no need to fear). In order to explain an event in terms of known mechanisms, they tend to ask whether the event meets the preconditions for the mechanisms.

Very few questions were directed as covariational relationships between the terms of the event (e.g., directly between John, Dave's raccoon, and the occasion). Those questions that did ask for covariation information seem to have been prompted by the repetitive wording of "on this occasion" in the no-schema sentences (i.e., "what about other occasions?").

Our manipulations in varied-schema sentences had very little overall effect. Even when presented with nonsense sentences ("The fep mimbled the wug") people attempted to apply their existing knowledge of causal relationships (e.g., "Is the fep mad at the wug?"). The tendency to search for further information about the target event held even in those cases where subjects might plausibly assume that they knew all the relevant facts (i.e., the "obvious" sentences: "The customer left the restaurant after eating"). We were unable to induce subjects to focus directly on the factors presented in the events and refrain from seeking more specific causal information. Finally, instructions that were aimed at biasing participants toward ANOVA questions and explanations were completely unsuccessful.

Explanations generated in this study almost invariably identified a mechanism for the relationship between cause and effect. If something was left ambiguous in subjects' explanations, it tended to be the particular factor responsible rather than the mechanism. For example, given a nonsense sentence "XB12 mimbled the wug at fulmer", a subject stated the event occurred because "The XB12 resented the wug". In this explanation, it is not clear whether it is due to something particular about XB12, the wug, or the occasion. Also, in many cases subjects were content to identify which factor was *not* the cause. For example, a subject said Dave would not eat rabbit meat because no one else did. This explanation looks like an ANOVA explanation (and the coders treated it as an ANOVA type using a lenient criterion). But it probably is not a true ANOVA explanation because we are still left to choose between the meat and the occasion (or their interaction) as the responsible factor. Rather the subject seems to be implying that social pressure caused Dave not to eat.

EXPERIMENT 2

The data from Experiment 1 show that people search for information and provide explanations in a way much more consistent with the predictions of the mechanism approach. However, the conclusions we can draw from Experiment 1 may be limited because subjects were asked to generate questions but were not provided with any answers. Although it is unclear why this should induce a bias for any particular kind of question or kind of explanation, it does introduce a substantial difference between our task and the more usual task of causal inference. In Experiment 2 we address this concern by providing answers to subjects' questions.

10. Method

10.1. Procedure

The experiment was run in groups of 1 to 3 subjects. Because of the length of the task, only half of the items from Experiment 1 were used. Half of subjects received eight no-schema sentences, and the other half received eight varied-schema sentences (marked with "*" in Appendices A and B). The procedure was similar to Experiment 1 except that only yes/no questions were allowed because, otherwise, it is extremely difficult to prepare answers for all possible questions.[7]

After each question, the experimenter wrote down the answer right next to the question according to a pre-determined answer list. The answer list was designed so that two thirds of the answers would be "yes" and one third "no". We used more yes than no answers because otherwise it might be too difficult for the subjects to come up with explanations if they asked a related line of hypothesis-testing questions as they did in Experiment 1. Sometimes the predetermined answer was inconsistent with an earlier answer. In that event, the predetermined answer was changed to be consistent with earlier answers. For example, given a target event "The customer left the restaurant after eating", if a subject asked "Did she have somewhere to go?" and the answer was "yes" (implying that the customer was female), then if the subject later asked, "Is the customer a man?", the answer had to be "no". During the experiment, no subject noticed that the answers were randomly provided by the experimenter.

Finally, the subjects were asked to write down an explanation for each event when they had nothing to ask about a given event.

[7] This constraint does not seem to be overly restrictive given that 74.4% of questions produced in Experiment 1 were yes/no questions and, furthermore, 87.0% of ANOVA questions were yes/no questions. Therefore, even if only yes/no questions were allowed in Experiment 2, it is unlikely that this task constraint would serve to reduce the number of ANOVA questions.

10.2. Subjects

There were 22 subjects who were undergraduate students at the University of Michigan, participating in partial fulfillment of course requirements for introductory psychology. Subjects were randomly assigned to one of the four conditions.

10.3. Scoring

Subjects' answers were classified into ANOVA questions, hypothesis-testing questions, and information-seeking questions. Very few questions were nonsense, linguistic, or effect questions. Nonsense questions (less than 1% of all the questions) were excluded from data analysis and linguistic and effect questions were treated as information-seeking questions.

11. Results

11.1. Questions

Only a small percentage of the questions were ANOVA questions (0.98% for the varied-schema sentences and 12.0% for the no-schema sentences; see Fig. 4).

The average number of each question type per item was used for an analysis of variance with events as a between-subject variable and question type as a within-subject variable. There was no reliable main effect of events, $p > .10$, a reliable main effect of question type, $F(2, 40) = 15.208$, $MSE = 1.076$, $p < .001$, and no reliable interaction effect of events by question type, $p > .10$. Pairwise comparisons of question type (Tukey test) showed reliably fewer ANOVA questions than either information-seeking questions or hypothesis-testing questions, $p < .01$.

For each of the 16 items, there were fewer ANOVA questions than hypothesis-testing questions. In addition, no subject produced more ANOVA questions than hypothesis-testing questions.

11.2. Analysis of order of questions

We examined whether the percentage of ANOVA questions varied depending on position in the sequence of questions. Since there were only three ANOVA questions produced in response to the varied-schema sentences, only the questions responding to the no-schema sentences were further analyzed. Fig. 5 shows the percentages of each question type in each position (i.e., absolute position in the sequence of questions) up to the fifth question because the pattern stabilizes afterward. As shown in the figure,

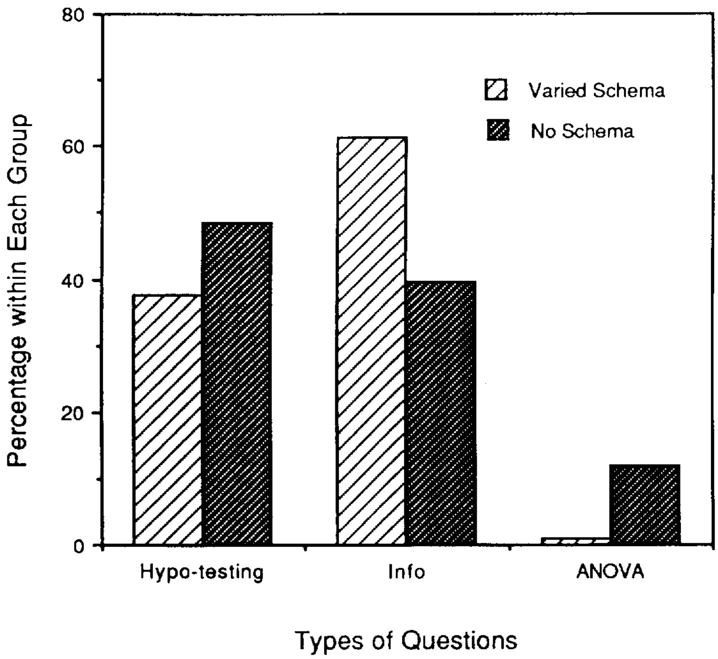

Fig. 4. Percentage of each question type for the varied-schema condition and the no-schema condition in Experiment 2.

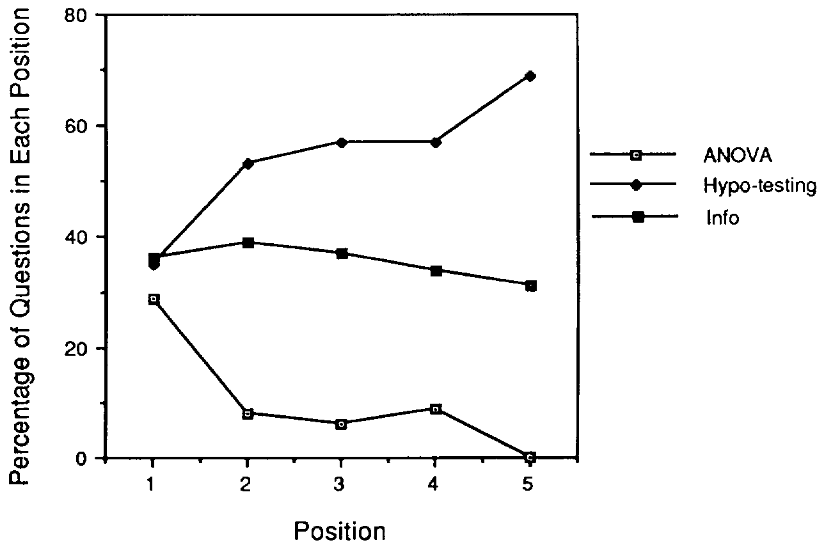

Fig. 5. Percentage of each question type in each position in Experiment 2.

what few ANOVA questions were asked appeared early in the sequence. Hypothesis-testing questions were more prominent and increased across sequences.

The following is an example of a sequence of questions where an ANOVA question was asked first (the information written in bold was given by the experimenter):

> Target event: **George likes collard greens (a leafy vegetable) on this occasion.**
> Does George eat collards often? **No**
> Can George cook collards himself? **Yes**
> Did George prepare them for someone else? **Yes**
> Is this a date for George? **Yes**
> Did George's date enjoy the collards? **No**
> Did George enjoy his entire date? **Yes**
> Is this their first date? **Yes**
> Is George in love? **Yes**
> Is she in love? **No**
> Explanation: George enjoys his cooking no matter what because he's used to eating alone anyway because he's a boring date.

In this example, the subject first asked an ANOVA question about the occasion and then kept on asking more hypothesis-testing questions to find out what is so special about the occasion. It is worth pointing out that even when she obtained information consistent with an occasion interpretation from the first (ANOVA) question, her explanation turned out to be an explanation not about occasion but rather about George because of the information she received from her later hypothesis-testing questions.

11.3. Analysis of explanations

From the varied-schema sentences, there were no ANOVA type explanations. From the no-schema sentences, there were only five ANOVA explanations. (Of these five, four were by the same subject.) The rest of the explanations were of the mechanism type.

12. Discussion

Results from Experiment 2 mirror those from Experiment 1. Hypothesis-testing questions dominated over ANOVA questions and the same pattern held for explanations. Subjects' questions suggest that they are looking for specific features of events which might indicate the presence of a familiar causal relationship. We failed to find any substantial evidence of covariational reasoning.

We found that ANOVA questions tended to be asked at the beginning of the sequence of questioning. Although ANOVA questions were hardly frequent enough that we should try to interpret this interaction, we offer one speculation. In some cases, one may fail to come up with a manageable set of hypotheses possibly because the search space is too large. Then one good heuristic might be to locate the source first. The answers to such covariation questions as "Was last night in any way special?" can be used to come up with a mechanism-based hypothesis (e.g., "The accident was due to a snow storm").

EXPERIMENT 3

Some researchers have suggested that the covariational approach to causal reasoning only characterizes people's behavior when they have all the covariation information needed to make a judgement (see Kelley, 1973). A similar argument is that people do not actively search for covariation information because such information is not typically available (Lipe, 1991). It seems to us that if covariation information is not typically available but is important in causal attribution, people should be rather good at looking for it. Although it is not clear to us why availability of covariation information might affect which information is important to causal attribution, Experiment 3 tested any possible effect of given covariation information because only a few studies have used information-seeking paradigm.

Subjects received varying amounts of covariation information and were instructed to ask more information or make causal attributions. This task allows us to address two questions. First, given partial information about covariation, what further information will participants *seek*? Second, will subjects *use* given covariation information as the basis of their attributions? As pointed out earlier, a large number of studies testing the *use* of covariation information have already shown results consistent with the ANOVA approach. The critical difference, however, between those studies and Experiment 3 lies in how open-ended the task is: subjects in the current experiment are free to give any type of explanations whereas previous studies allowed subjects to choose only among ANOVA or covariational explanations.

13. Method

13.1. Design and materials

All items were taken from Cheng and Novick (1990a). Each event description varied in the amount of covariation information. In the no-information condition, there was a single target event description (e.g., "Eric did not like to dance the tango on this occasion"). In the partial-information condition, there was a single target event description plus three pieces of covariation information. Depending on how the three pieces of covariation information were selected, there were two groups. In the systematic group, the partial covariation information was about consistency, consensus, and distinctiveness (e.g., "Nobody else liked to dance the tango on this occasion. Eric did not like to dance any other kind of dance on this occasion. On all other occasions, Eric has never liked to dance the tango"). In the random group, the three pieces of covariation information were randomly selected from the seven pieces of covariation information specified in Fig. 1. Five combinations of randomly selected covariation information were used (2, 5, and 8; 3, 6, and 7; 4, 6, and 8; 2, 3, and 6; 2, 4, and 7 in Fig. 1). In the complete-information condition, there were a total of eight statements (i.e., cells 1 through 8 in Fig. 1) including a target event description and complete covariation information presented. (See Appendix C for target event descriptions plus complete covariation information descriptions for all events.)

Each participant received 15 event descriptions, five of which were in the no-information condition, five in the partial-information condition, and five in the complete-information condition. Three randomized combinations were used to determine which events were to be in each condition. Half of the subjects received partial covariation information randomly selected from the complete covariation information set and the other half received consistency, consensus, and distinctiveness as partial covariation information.

13.2. Procedure

Subjects received a booklet containing instructions and event descriptions. The subjects were asked to generate questions and explanations for each event, but this time the instructions were focused on producing explanations rather than questions as follows:

> Each event description consists of two parts. The first sentence in each problem is a target event that you are to explain. This sentence is underlined. In some of the problems, we added some more information that may help you to determine the cause of the target event. However, if you would like to ask for more information concerning the event, write the questions with a pen other than a red pen.

Subjects carried out their task without receiving any given answers to their questions.

13.3. Subjects

There were 29 subjects who were undergraduate students at the University of Michigan, participating in partial fulfillment of course requirements for introductory psychology. Fourteen subjects were randomly assigned to the systematic group and 15 to the random group.

14. Results

14.1. Overall results of questions

Using the same scoring method as in Experiment 2, we found that 42.7% of the subjects' questions were of the ANOVA type, only 27.3% were of the hypothesis-testing type, and 30% were of information-seeking type. Fig. 6 shows the frequencies of each type of question for each group in each condition. In both the random and systematic groups, there was no reliable difference in frequency of each type of question, $p > .10$. Compared to the first two experiments the proportion of ANOVA questions have dramatically increased but the absolute frequency was small averaging 0.44 ANOVA questions per item. Furthermore, the ANOVA questions were produced mainly by a subset of subjects: only three out of the total of 29 subjects produced more ANOVA questions than non-ANOVA questions. Among these three subjects, one subject by the end of the experiment completely figured out the seven pieces of covariation information and showed complete covariation reasoning. In addition, 49% of the ANOVA questions were produced by only five subjects.

14.2. Analysis of ANOVA questions

An analysis of variance was carried out to test whether the number of ANOVA questions varied depending on group (random or systematic) and amount of information. There was no reliable main effect of group, $p = .08$ and no reliable interaction effect of group and condition, $p = .09$. The main effect of amount of information was reliable, $F(2, 54) = 17.784$, MSE = 10.756, $p = .0001$. Subjects tended to ask more ANOVA questions in the no-information condition (1.01 questions per item, or 53.6% of the questions asked in this condition) than in the partial condition (0.29 questions per item, or 42% of the questions asked in this condition) and the complete condition (0.02 questions per item, or 3.9% of the questions asked in this condition). Still, the average frequencies of ANOVA questions were not even close to the expected numbers: if subjects showed complete covariational reasoning in the no-information condition, then there should have been seven ANOVA questions, yet the average frequency was only 1.01. Similarly, in the partial condition, the expected number of ANOVA

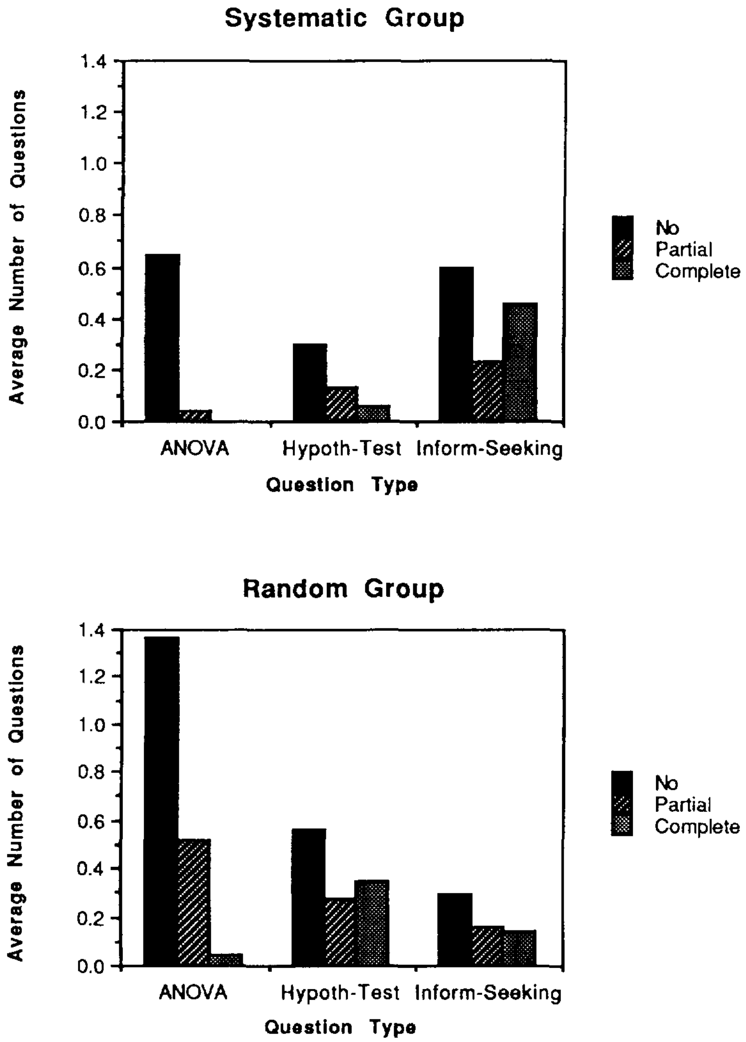

Fig. 6. Percentage of each question type for each information condition for the systematic group and the random group in Experiment 3.

questions was 4 (because three pieces of covariation information out of seven are given in the problems), yet the average frequency was only 0.29.

14.3. Overall results of explanations

On average, only 17% of the explanations were of ANOVA type and the rest were either mechanism explanations (73.1%) or nonsense explanations (9.9%). Most importantly, even from the complete covariation condition, ANOVA explanations were given only 22.7% of the time (31.5% by the random group and 13.8% by the systematic group). For the systematic group, the partial information condition actually provided complete covariation information in Kelley's sense, yet only 12.3% of these explanations were ANOVA explanations. Table 2 shows the proportion of each explanation type for each group broken down into the amount of covariation information.

14.4. Consistency of explanations

Both ANOVA and mechanism explanations were coded for types of attribution (person, occasion, stimulus or their combinations). Cheng and Novick's probabilistic contrast model is most readily applied to the complete covariation condition. The predictions were directly taken from Cheng and Novick (1990a) and both weak and strong causes were considered. Of the total of 145 explanations of any type (i.e., both ANOVA and mechanism explanations), 39.3% were one of the strong causes predicted by the probabilistic contrast model. When weak causes were also considered, 62.1% were one of the causes predicted by the probabilistic contrast model. Of the total of 31 ANOVA explanations, only 12 explanations (or 39%) were consistent with the strong causes and five more explanations were consistent with the weak causes.

To provide a baseline to evaluate this success, we compared subjects'

Table 2
Proportion of each explanation type for each group broken down into the amount of covariation information in Experiment 3

Covariation	Explanation Type			
	ANOVA	Hypothesis-testing	Nonsense	Total[a]
Systematic group				
None	0.005	0.31	0.02	0.34
Partial	0.04	0.25	0.04	0.33
Complete	0.05	0.26	0.03	0.33
Total	0.09	0.83	0.08	
Random group				
None	0.07	0.23	0.01	0.31
Partial	0.07	0.24	0.04	0.35
Complete	0.11	0.18	0.06	0.34
Total	0.24	0.64	0.11	

[a] The total for each row is not exactly 0.33 because some subjects did not come up with explanations.

attributions in the no-information condition to the probabilistic contrast model's predictions for the same event. The predictions require subjects to have complete covariation information. Since no covariation information was provided in the no-information condition, there is no reason to expect accurate predictions for these events. Thus the success rate of the probabilistic contrast model's predictions in the no-information condition can be taken as a measure of chance responding. Interestingly, 47.0% of subjects' explanations in the no-information condition were consistent with predictions made by the probabilistic contrast model (both strong and weak causes). Therefore, the complete-information condition did not seem to have much advantage over the no-information condition.

The data were also compared with predictions made by Kelley's ANOVA model. Kelley's model can make predictions in the complete-information conditions and the partial-information condition for the systematic group. (Subjects in these conditions received information on consistency, consensus, and distinctiveness.) Only Items 1, 2, 3, 4, 5, and 6 in Appendix C were considered, because the ANOVA model does not make any predictions about the other items. Across all types of explanations (both mechanism and ANOVA), 35.6% of the explanations were consistent with the ANOVA model's predictions. Among ANOVA explanations (12 in total for these items in the conditions specified above), only five were consistent with the ANOVA model's predictions. Given no covariation information, 15.7% of subjects' explanations were consistent with the ANOVA model's predictions.

15. Discussion

On the one hand, we were able to increase proportions of covariation questions and such requests did seem sensitive to the amount of presented covariation information (i.e., there was an inverse relation as predicted by covariational models). Interestingly, however, there were twice as many ANOVA questions in the no-information condition (1.01 per item) than in the first two studies (0.51 and 0.48, respectively) when the materials were the same across these studies. Clearly, the presentation of covariation information to participants was associated with their asking for more covariation information. Note also, however, that about half of the total number of ANOVA questions were in fact produced by only 17% of the subjects. Among these subjects, one figured out the seven pieces of covariation information after solving about half of the problems. One interpretation of this finding is that it reflects a response to demand characteristics – "the experimenters seem to think these kinds of questions are good so I'd better ask some". An alternative view is that participants may not normally seek out covariation information but they recognize its value and realizing its effectiveness, begin to request it. Although the present study cannot differentiate these two possibilities, the analysis of subjects clearly shows that ANOVA questions were produced by the minority of the subjects.

This experiment also indicates that the request for covariation information is not the same as the conduct of an analysis of covariation. When subjects did ask for covariation information they did not ask for enough to conduct a complete analysis. Typically only one or two covariation questions would be asked per event. Usually some hypothesis-testing questions were asked together with the covariation questions. In addition, as in the first two studies, the overwhelming majority of explanations referred to some mechanism for the effect. ANOVA explanations were uncommon and when they did occur, they did not match the predictions of the covariation models.

EXPERIMENT 4

The three experiments reported so far suggest that covariation information is not centrally the most salient aspect of people's causal reasoning. Advocates of covariational models might suggest that the data we have presented so far merely demonstrate that people do not ask for information or produce explanations in the form we recognize as covariation. Although subjects are producing hypothesis-testing questions, they might be primarily interested in the covariation information such a question would provide. Any statement reporting the mechanism whereby a factor might have acted as a cause will also *indirectly* report a covariation between that factor and the cause. For instance, stating that "John does not know how to drive" could be interpreted as a statement about how "John" covaries with "accidents". According to this line of reasoning, subjects might appear to be focusing on the mechanisms present in events, but the effective information would be the covariation information indirectly conveyed by mechanism information.

Experiment 4 evaluates this idea. Is there something over and above information about covariation that is conveyed by a statement of a mechanism for a relation? Is the information that John does not know how to drive identical to the information that John is particularly likely to have car accidents? We test this issue in the following way: one group of subjects is asked to assess the covariation values conveyed by mechanism statements. We then give other subjects different combinations of mechanism and direct covariation statements and assess the strength of their causal attributions. If mechanism statements only function by means of providing covariation information, then we should be able to use the first group's covariation ratings to predict causal attributions made by mechanism statements.

16. Method

16.1. Materials

We developed nine event descriptions including person, stimulus, and occasion components. (See Appendix D for a complete list.) For each event, two questions were asked in the form of "To what degree was this event due to X?" For three events, X was the person in one question and the stimulus in the other; for another three events, it was the person and the occasion; and for another three events, it was the stimulus and the occasion. For each factor in question, we developed an explanation conveying mechanism information. In Appendix D, these sentences are followed by "Mechanism (the name of the corresponding factor)".

16.2. Assessing covariation value of mechanism statements

For each mechanism sentence, we developed a corresponding covariation sentence through the following separate experiment with 28 subjects. In this pretest, the subjects received a list of the mechanism sentences and were asked, for each sentence, to rate how much more or less likely the situation described in the mechanism sentence made the corresponding target event. For example, given that "There was a very severe storm and the roads were very slick last night", the subjects were asked "Compared to the average night, how likely were people to have a traffic accident last night?" The subjects answered these questions by rating on a 13-point scale marked with extremely less likely (−6), very much less likely (−5), much less likely (−4), somewhat less likely (−3), a bit less likely (−2), barely less likely (−1), about as likely (0), barely more likely (1), a bit more likely (2), somewhat more likely (3), much more likely (4), very much more likely (5), and extremely more likely (6). Covariation sentences were created using the phrase corresponding to the mean ratings. In the example above, if the mean rating was closest to "much more likely", then the covariation sentence was "Traffic accidents were much more likely last night than on other nights". The covariation sentences are presented in Appendix D followed by "Covariation (the name of the corresponding factor)".

16.3. Design

There were six conditions depending on which combination of mechanism and/or covariation information subjects received for each factor in a given target event. Table 3 shows our notation for the conditions, which combination of information the subjects received, and which sentence combinations in Appendix D correspond to which condition.

Table 3
Summary of conditions in Experiment 4

Name of condition	Combination of additional information	Sentence numbers in Appendix D
None	None	None
C/C	Two covariation sentences, one for each factor	2 and 4
M/C	Mechanism sentence for one factor and covariation sentence for the other factor	1 and 4, or 2 and 3
M/M	Two mechanism sentences, for one each factor	1 and 3
MC/M	Mechanism and covariation sentences for one factor and mechanism sentence for the other factor	1, 2, and 3, or 1, 3, and 4
MC/C	Mechanism and covariation sentences for one factor and covariation sentence for the other factor	1, 2, and 4, or 2, 3, and 4
MC/MC	Mechanism and covariation sentences for both factors	1, 2, 3, and 4

16.4. Procedure

The subjects received the following instructions.

In this experiment, you will receive descriptions of events (e.g. Harry punched Jack). These target descriptions are underlined. Below each event description, you will see additional statements related to the event (e.g. "Harry is a violent person". "Jack is much more likely to be punched by other people".). Then, we'll ask you to rate the degree to which each factor was responsible. For example, you will see questions as follows:

Harry punched Jack.
Harry is a violent person.
Jack is much more likely than other people to get punched.

To what degree was this event due to something particular about Jack?

Mostly due to Jack   Not due to Jack

1   2   3   4   5   6   7

To what degree was this event due to something particular about Harry?

Mostly due to Harry   Not due to Harry

1   2   3   4   5   6   7

Subjects also received explanations for such terms as "much more likely", "extremely less likely" by receiving the same scale used in the pretest and they were told to use the scale when they solved the problems. The order of the nine problem sets was completely randomized across all subjects. The questions in each problem set and the additional sentences (mechanism and/or covariation sentences) were randomized in two different orders.

16.5. Subjects

The subjects consisted of 197 undergraduate students at the University of Michigan, either participating in partial fulfillment of course requirements for introductory psychology or paid $5.00 an hour for participating in this experiment. All subjects were randomly assigned to one of the conditions.

17. Results

In this experiment we are interested in how two variables affect a factor's rating of causal strength. The first variable involves the kind of information supporting the factor (i.e., mechanism, covariation, or both kinds). The second variable involves the information supporting a competing factor (again, mechanism, covariation or both). Mean responses for different levels of these variables are shown in Table 4. The smaller the number, the stronger the causal attribution.

Mechanism information was more influential than covariation. For example, when mechanism information for a factor was competing against mechanism information for another factor, attributions for the factor were weaker (3.43) than when mechanism information was competing against covariation information for the other factor (2.74). Similarly, when both

Table 4
Mean ratings for each condition in Experiment 4

Information for	Information against	Average rating
Mechanism	Mechanism	3.43
Mechanism	Covariation	2.74
Mechanism	Both	3.68
Covariation	Mechanism	3.75
Covariation	Covariation	3.29
Covariation	Both	4.04
Both	Mechanism	2.91
Both	Covariation	2.55
Both	Both	3.28
None	None	3.50

mechanism and covariation information was competing against mechanism information (2.91), the attributions were weaker than when they were competing against covariation information (2.55). When covariation information was competing against mechanism information (3.75), the attribution was also weaker than when it was competing against covariation information (3.29). But covariation information was not completely ineffective: compare mechanism information competing with mechanism versus mechanism competing with both mechanism and covariation information (3.43 vs. 3.68). As the following more formal analysis shows, mechanism information was roughly twice as effective as covariation information.

17.1. Direct competition model

We assume that each kind of information (mechanism and covariation) has a particular weight on attributions. When the information supports the attribution, the weight is positive; when the information supports some other factor, the weight is negative. A subject's attribution would be computed using the following equation:

$$\begin{aligned}\text{Attribution} = K &+ \beta \text{ (mechanism for)} - \beta \text{ (mechanism against)} \\ &+ \partial \text{ (covariation for)} - \partial \text{ (covariation against)}\end{aligned}$$

The mechanism weight (β) will be larger than the covariation weight (∂) to the extent that mechanism information contributes more than implied covariation.

17.2. Testing the model

Since subjects made two ratings for each event (e.g., rating the power of the person factor and the power of the occasion factor), a single piece of information served simultaneously as a support and as a competitor. For instance, the piece of information "Kim does not wear her glasses" supports the person factor and competes with the occasion factor. Thus, whether something is a support or a competitor depends on the factor ratings we are looking at.[8] We have split the data and performed two separate analyses. In every case, the two factors that were rated for a single event will appear in

[8] To analyze both effects (of support and competition) we could do a single analysis with which factor was being rated as a within-subjects variable. This strategy would mean, however, that no main effects would be significant; the effect of a piece of information depends on which factor rating we examine.

separate analyses. The information that is supporting in one case will be competing in the other.

17.3. Regression analysis

The direct competition model can be tested in a regression analysis. The model contains two predictor variables: mechanism and covariation. The variables have the value of 1 when that particular kind of information supports the attribution and a value of -1 when the information competes. The variable takes the value of 0 when the information both supports and competes (i.e., the hypothesis is that the effect of the information cancels out). For example, a subject in the MC/C condition would be represented by a mechanism variable of 1 (supporting mechanism) and a covariation variable of 0 (canceling out of effects). Data were entered into one of two, identical regressions (regression 1 and regression 2) with mechanism $(1, 0, -1)$ and covariation $(1, 0, -1)$ as predictors.[9]

Both regressions accounted for a significant portion of the variance (regression 1: R-squared = .30, $F(2, 173) = 37.47$, MSE = .55, $p < .001$; regression 2: R-squared = .25, $F(2, 173)$, MSE = .60, $p < .001$). Although the R-squares seem low, they are significant given that we are only using 2 degrees of freedom to predict 175 data points. Beta weights from the two regressions are shown in Table 5. The confidence intervals for the weights of mechanism and covariation information do not overlap ($p < .05$) in either regression. The magnitudes of the mechanism weights are close to double those of the covariation weights. This indicates that mechanism information has about twice as much impact on attributions. (See Appendix E for a test of a different model.)

Table 5
Beta weights for two regression analyses in Experiment 4

Variable	Beta weight	90% upper	90% lower	*P* value (weight differs from 0)
Regression 1				
Mechanism	−.83	−.99	−.67	.0001
Covariation	−.36	−.52	−.20	.0002
Regression 2				
Mechanism	−.75	−.92	−.59	.0001
Covariation	−.30	−.47	−.13	.003

[9] Data from the no-information condition were not included in the analyses. Subjects in this condition seem to be performing a very different task from subjects in the other conditions, who had at least some supporting and competing information given. Further, in the model we are fitting, this condition is represented as 0,0 – a configuration of predictor variables already oversampled (by data from M/M, C/C, and MC/MC).

18. Discussion

The results showed that subjects' attributions are more affected by mechanism statements than covariation statements even when the mechanism statements were empirically determined to convey the same covariation values. These results demonstrate that there is more information conveyed by a mechanism than just covariation.

One might argue that the mechanism sentences and the covariation sentences are different not only with respect to the presence or absence of mechanism information but also with respect to other factors. For example, one might say that covariation-based explanations used in the experiment sound too unnatural or too formal to be used in natural conversation and that might be why they were less effective. We argue that the criticism actually supports our claim; the reason the covariation-based explanations sound weird is because people usually do not talk this way. Most of explanations that we use in everyday situation sound natural because they convey mechanism information. The only way we could present bare covariation information minimizing the possibility of suggesting any mechanism in our experiments is to use such a formal expression.

In addition, the mechanism sentences might look more specific and more informative than the covariation sentences, making them more *reliable* covariation information than the covariation sentences. Note, however, that all the mechanism sentences are not necessarily reliable predictors of the target events because, for example, a person not wearing her/his glasses might or might not have traffic accidents (if they compensated by driving more cautiously they might even have fewer accidents). In the current experiment, the covariation sentences were developed by measuring the subjects' pre-existing background knowledge about this covariation. Therefore, as far as the amount of covariation information is concerned, there were no differences between the mechanism sentences and the covariation sentences.

GENERAL DISCUSSION

19. Summary and implications of results

The four experiments presented here suggest that people do not spontaneously seek out information about covariation between factors and effects, nor do they seem to use such information when it is provided. The preferred strategy seems to be to gather further facts about the particular event to be explained: facts that are used to test hypotheses about possible underlying mechanisms.

In Experiments 1 and 2, subjects were asked to indicate the questions they would like answered in order to discover the causes of some events. In this task, people sought out information about target events rather than seeking information about patterns of covariation between factors and effects in other events. Questions reflected hypotheses about mechanisms ("Was John drunk?") rather than general strategies of covariation analysis ("Has John had other accidents?"). This pattern held across a variety of event types including nonsense sentences ("The fep mimbled the wug") and the no-schema sentences where there should not have been obvious or familiar hypotheses about possible causes. There seems, then, to be a strong bias away from seeking out information about patterns of covariation and a tendency to focus on specific features of target events that would provide evidence for or against hypotheses about causal mechanism. The results from nonsense sentences and the no-schema sentences suggest that this bias or tendency is not simply an effect of previously learned relationships.

In Experiment 3, covariation information was given in some of the problems (partial and complete covariation conditions) but not in others (no-covariation condition). Subjects' task was, again, to indicate any further questions necessary for causal attribution. In this experiment, more questions about covariation were produced than about hypothesized mechanisms, but much of this result was due to a few subjects who overwhelmingly asked covariation questions. Most subjects typically requested only one or two pieces of covariation information (which are insufficient to identify a covarying factor) and more frequently asked questions about mechanisms. It is at least plausible that the increased proportion of covariation questions resulted from subjects modeling their questions on the partial covariation information that was provided in the experiment. Also, in Experiment 3, the kinds of explanations subjects gave provide evidence that the subjects were not performing a covariation analysis. In a subset of the cases from Experiment 3, subjects were provided with sufficient information to conduct a covariation analysis. Subjects' attributions in these cases were tested against the predictions of two covariation models (Kelley, 1967; Cheng & Novick, 1990a, 1990b, 1991, 1992). No effect of covariation information was found. Our subjects' attributions accorded with the models to the same degree whether or not they were provided with any covariation information. The majority of explanations stated some mechanism not present in the statement of the target event.

Finally, Experiment 4 investigated whether the use of mechanism information involves more than the use of covariation information implied in the mechanism. One alternative interpretation of Experiments 1, 2, and 3 is that when asking non-ANOVA questions subjects might have only been interested in the information about covariation implicitly conveyed by information about mechanisms. In Experiment 4, subjects were presented with pieces of conflicting evidence matched for degree of covariation information provided. We varied whether the information included some statement of a mechanism for the relationship stated. The effect of evidence was stronger when the evidence included mechanism information. This effect is over and above the effect of covariation information included in the mechanism statements.

20. Relationship between mechanism and covariation information

Although we failed to show prevalent use of or search for covariation information, this should not be taken as an argument that such information is never important in identifying causes of events. The process and the purpose of identifying causes is clearly very complex and may involve a complicated relationship between covariation and mechanism information (and probably other factors that were not considered in the present study). Indeed, the relationship between covariation and mechanism seems very much like, and is importantly related to, the relationship between similarity and theory in categorization (Medin, 1989; Gelman & Medin, 1993; Murphy & Medin, 1985). This last section presents our ideas about how covariation and mechanism information might each be useful under different circumstances.

20.1. Mechanism versus covariation and similarity versus theory

Models of similarity-based concept formation and covariation models of causal attribution describe the same kind of process. In similarity-based categorization, instances are described in terms of a set of independent features, just as the covariation-based models represent events in terms of a set of factors. Roughly speaking, similarity-based concept formation proceeds by abstracting common features and discarding different features across members in the same category. A causal identification based on covarariation analysis fits this framework. Suppose the target event is "John had a car accident last night" and the question is, for example, whether John was the cause of the accident. The next step is to find out whether John was present in other similar cases (e.g., cases of John driving other cars, cases of other people driving on similar roads). If no factor other than John was present in all these cases, then John would be identified as the cause of the event.

In contrast, in cases of mechanism-based causal attribution and theory-based concept formation, the goal is to identify instances of laws and processes at work in the target case. In theory-based categorization, categories are coherent not because of common features but because of common explanatory structures (Ahn, Brewer, & Mooney, 1992; Medin, 1989; Murphy & Medin, 1985). For example, pets, children, photo albums, and diaries are not perceptually similar to each other, but they belong to a category of "things to take out from one's house in case of a fire" because of our theories about the objects and the category (Barsalou, 1983). Mechanism-based causal attribution also proceeds by finding some principle or rule which could account for the current case. For example, the discovery that John was asleep at the wheel allows us to identify that as the cause of the accident, not because we know that sleeping is correlated with accidents, but because we understand processes that could lead to accidents if the driver were asleep.

Given this parallel between covariation and similarity, and mechanism and theory, it becomes clear that the relationship between covariation and mechanism will be complex, because in categorization both similarity and theory play important roles depending on contexts. As many researchers have argued (Gelman & Medin, 1993), concepts are multidimensional, having various functions (e.g., categorization, inference, organizing information, problem solving). Because similarity and theory play different roles depending on the function, it would be meaningless to ask which one is more important. We would argue that processes involving covariation and mechanism stand in exactly the same relationship. In the following section, we will briefly consider some of the rationales and motives for using both covariation and mechanism information in various contexts of identifying the causes of events.

20.2. *Covariation as a heuristic for hypothesis generation*

Mill (1843/1973) makes the distinction between contexts of discovery, roughly hypothesis generation, and contexts of justification, roughly hypothesis testing. Covariation methods would be quite useful in generating hypotheses. The methods can efficiently sort out irrelevant factors and reduce the size of the hypotheses to be tested. In addition, when one has difficulty coming up with a mechanism-based hypothesis, one might first attempt to identify a responsible factor. In our Experiment 2, the subjects tended to ask covariation questions at the beginning of the sequences of questioning. This result seems to suggest that covariation information would be useful when people have difficulty generating mechanism-based hypotheses.

20.3. *Covariation as evidence and mechanism as theory*

A second way covariation information might be useful is in justifying the presence or absence of a particular mechanism. That is, the relationship between covariation and mechanism is one of evidence and theory. When searching for the mechanism that caused an event, one might gather information about covariation rather than, or in addition to, direct information about the mechanism. This might be a particularly useful strategy when direct information about mechanisms is difficult or impossible to attain. For example, it is typically hard to gather direct evidence about the processes involved in human actions (e.g., Did John steal the wallet because he is greedy or because he wants attention?). A direct inspection of the mental process that went into the event is not feasible. Instead we might inspect John's behavior in other situations (e.g., Does he generally show a lust for money?). When information about mechanisms are not readily available, covariation methods can be useful heuristics because covariation is perfectly confounded with mechanism: the presence of a mechanism necessarily implies the presence of covariation among the factors involved in the mechanism. To draw another analogy between categorization and causal attribution, superficial similarities (e.g., hair length, height for "boy" concept) often serves as an indicator of underlying explanatory structures (e.g., male chromosomes for "boy" concept) (Medin & Ortony, 1988; Murphy & Medin, 1985). Depending on the resources available, covariation information might be an excellent source of data to use in justifying hypotheses about mechanisms. We can find this type of strategy in real life situations; most epidemiological studies are not directly concerned with mechanism. The finding that consumption of high levels of salt covaries with hypertension may be useful and important even when we do not know the mechanism whereby salt causes high blood pressure.

20.4. Conversational pragmatics

Covariation might play a central role because of conversational pragmatics (Hilton, 1990). An explainer simply needs to refer to the source of the cause when the listener shares the same mechanism knowledge. For example, one might say a person didn't go to work because it was Sunday (i.e., something special about the occasion). In this case, just the statement of a responsible factor will often be enough for us to complete the story of how the event happened.

20.5. Role of explanations

Why might people prefer the mechanism-based explanations over the covariation-based explanations? The answer depends on the role of each type of explanation. Simply identifying responsible factors at the surface level (e.g., something about John as a cause of his traffic accident) does not give an attributer the evaluative component for attribution. John could have been valiantly rushing rushing an injured person to the hospital or he could have been drunk. On the other hand, mechanism-based explanations are projective or generative in the sense that we can make predictions about novel situations. As discussed earlier, the mechanism approach argues that the causal relationship is one of necessity. This characteristic of causal necessity allows an explainer to project the causal relationship to novel situations. This is often talked about in terms of counterfactuals (Hilton & Slugoski, 1986; Lipe, 1991). Suppose John had a traffic accident and the mechanism-based explanation for the event is that John was drunk. What if a particular factor had been different? For example, what if John was wearing a hat? Would the effect still have occurred? The explainer already knows the preconditions for accidents involving drunk driving and he/she knows that the effect would occur by necessity as long as the conditions are satisfied. Therefore, the projection into this new situation can be easily carried out. If we do or don't want the effect to occur again, what can we do? The explanations based on mechanisms or causal necessity also allow us to control the events in the future.

In contrast, the identification of a mechanism gives us some reason to believe that we can project the effect by manipulating the factor. For example, the observation that putting your hand near a fire has led to warmth in the past does not guarantee it will in the future. On this (Humean) analysis it is just habit that causes us to go to fires when we want to get warm. In contrast, if we understand something about the mechanisms of fire, heat, conductivity, etc., we have good reasons to believe that the fire will warm us. There is some necessity to the relationship between fire and warmth that assures us the relation will continue to hold in the future (Harré & Madden, 1975).

21. Implication

The current study presents a new way of understanding causal attribution processes. Unlike traditional approaches based on covariational analyses, the experimental results suggest that people are primarily interested in searching for mechanism information (i.e., how a factor caused the event). However, rather than competing, the two approaches have complementary roles. For example, there might be contexts in which people use covariation analysis to identify a causal factor and those in which they will seek to identify a mechanism for the effect. Mechanism information can be used to distinguish between true causality and spurious correlation from the product of covariational analysis and also has several fundamental implications such as projection and beliefs about conditions that might lead to change.

Acknowledgements

We would like to thank Hal Arkes, Patricia Cheng, John Coley, Evan Heit, Edward Wisniewski, and three anonymous reviewers for their helpful comments on the earlier drafts. This research was supported by NSF BNS 91-10245 given to Douglas L. Medin.

Appendix A: materials used for no-schema condition in Experiments 1 and 2

Jane had fun washing dishes on this occasion.
Cathy is not allergic to olive tree pollen on this occasion.
Adam thinks that narcissus (a flower) smells nice on this occasion.*

Wendy does not like to drink Campari (a kind of liquor) on this occasion.*
Sam is not afraid of John's raccoon on this occasion.
Beth said the morning prayer on this occasion.
Dave would not eat rabbit meat on this occasion.
Vicky is not wearing a gardenia (a flower) on her collar on this occasion.*
Patty enjoyed playing majong (a gambling game) on this occasion.*
Kim does not enjoy listening to the zither (a musical instrument) on this occasion.*
Alice displayed sculptures made from clay in her home on this occasion.*
Fred sang "Golden Slumbers" (a children's song) on this occasion.
Eric did not like to dance the tango on this occasion.*
Susan did not bow to the statute of fire on this occasion.
George likes collard greens (a leafy vegetable) on this occasion.*

Appendix B: materials used for varied-schema condition in Experiments 1 and 2

Bad events

Flight 921 crashed in Lincoln Nebraska last winter.*
The ballerina slipped and fell on the stage during opening night's performance.
The lawyer's houseplant died last week.
Harry killed Hanna on the way from work.*

Good events

The musician learned to play chess over the summer.
The student performed perfectly on the final exam in sociology.*
The athlete beat the city record in high jump during training.
Marry got a job promotion last year.*

Nonsense sentences

The pilson wizzled the globuleum during the vebel.
The toves gyred in the ipnoon.*
XB12 mimbled the wug at fulmer.*
A fep meckled a moja last Bom.

Obvious events

Susan closed the window when the rain started.*
The customer left the restaurant after eating.*
The professor lectured to the class on Tuesday.

The driver stopped his car at the intersection when the light turned red.

Appendix C: materials used for Experiment 3

Note: Although Cheng and Novick (1990a) used complete covariation information, each sentence contained different amounts of information with respect to the cube in Fig. 1. For example, a sentence could be about one cell in Fig. 1 or about two or three cells in the figure. In our materials, we rewrote all sentences in such a way that each sentence was about one cell. The content of the sentences remained the same as in Cheng and Novick.

Item 1

1. Jane had fun washing dishes on this occasion.
2. Nobody else had fun washing dishes on this occasion.
3. Jane han fun doing all chores on this occasion.
4. Jane always has fun washing dishes.
5. Nobody else ever had fun washing dishes.
6. Jane always has fun doing all chores.
7. Nobody else had fun doing all chores on this occasion.
8. Nobody else ever has fun doing all chores.

Item 2

1. Cathy is not allergic to olive tree pollen on this occasion.
2. Everybody else is allergic to olive tree pollen on this occasion.
3. Cathy is not allergic to any other kind of tree pollen on this occasion.
4. On all other occasions, Cathy is not allergic to olive tree pollen.
5. On all other occasions, nobody else is allergic to olive tree pollen.
6. On all other occasions, Cathy is not allergic to any other kind of tree pollen.
7. Everybody else is allergic to any other kind of tree pollen on this occasion.
8. On all other occasions, nobody else has been allergic to any other kind of tree pollen.

Item 3

1. Adam thinks that narcissus (a flower) smells nice on this occasion.
2. Nobody else thinks that narcissus smells nice on this occasion.
3. Adam thinks all other flowers smell nice on this occasion.
4. Adam has always thought that narcissus smells nice.
5. Nobody else has ever thought that narcissus smells nice.
6. Adam has always thought that all other flowers smell nice.

7. Everyone else thinks that all other flowers smell nice on this occasion.
8. Everyone else has always thought that all other flowers smell nice.

Item 4

1. Wendy does not like to drink Campari (a kind of liquor) on this occasion.
2. Nobody else likes to drink Campari on this occasion.
3. Wendy likes to drink all other kinds of alcoholic drinks on this occasion.
4. Wendy has never liked to drink Campari.
5. Nobody else has ever liked to drink Campari.
6. Wendy has always liked to drink all other kinds of alcoholic drinks.
7. Everyone else likes to drink all other kinds of alcoholic drinks on this occasion.
8. Everyone else has always liked to drink all other kinds of alcoholic drinks.

Item 5

1. Sam is not afraid of John's raccoon on this occasion.
2. Nobody else is afraid of John's raccoon on this occasion.
3. Sam is afraid of all other raccoons on this occasion.
4. Sam has never been afraid of John's raccoon.
5. Nobody else has ever been afraid of John's raccoon.
6. Sam has always been afraid of all other raccoons.
7. Nobody else is afraid of all other raccoons on this occasion.
8. Nobody else has ever been afraid of all other raccoons.

Item 6

1. Beth said the morning prayer on this occasion.
2. Everyone else said the morning prayer on this occasion.
3. Beth did not say any other prayer on this occasion.
4. On all other occasions, Beth always had said the morning prayer.
5. On all other occasions, everyone else said the morning prayer.
6. On all other occasions, Beth said all other prayers.
7. Nobody else said all other prayers on this occasion.
8. On all other occasions, everyone else said all the prayers.

Item 7

1. Dave would not eat rabbit meat on this occasion.
2. Nobody else would eat rabbit meat on this occasion.
3. Dave would not eat any kind of meat on this occasion.
4. On all other occasions, Dave has eaten rabbit meat.
5. On all other occasions, everyone else has eaten rabbit meat.

6. On all other occasions, Dave has eaten any kind of meat.
7. Nobody else has eaten any kind of meat on this occasion.
8. On all other occasions, everyone else has eaten all kinds of meat.

Item 8

1. Vicky is not wearing a gardenia (a flower) on her collar on this occasion.
2. Nobody else is wearing a gardenia on his or her collar on this occasion.
3. Vicky is not wearing any other kind of flower on her collar on this occasion.
4. On all other occasions, Vicky has worn a gardenia.
5. On all other occasions, everyone else has worn a gardenia on his or her collar.
6. On all other occasions, Vicky has never worn any other kind of flower.
7. Nobody else is wearing any other kind of flower on his or her collar on this occasion.
8. On all other occasions, nobody else has ever worn any other kind of flower on his or her collar.

Item 9

1. Patty enjoyed playing majong (a gambling game) on this occasion.
2. Everyone else enjoyed playing majong on this occasion.
3. Patty enjoyed playing all other kinds of gambling games on this occasion.
4. On all other occasions, Patty did not enjoy playing majong.
5. On all other occasions, everybody else enjoyed playing majong.
6. On all other occasions, Patty did not enjoy playing all other kinds of gambling games.
7. Everyone else enjoyed all other kinds of gambling games on this occasion.
8. On all other occasions, everyone else enjoyed playing all other kinds of gambling games.

Item 10

1. Kim does not enjoy listening to the zither (a musical instrument) on this occasion.
2. Nobody else enjoys listening to the zither on this occasion.
3. Kim does not enjoy listening to all other instruments on this occasion.
4. On all other occasions, Kim does not enjoy listening to the zither.
5. On all other occasions, nobody else has enjoyed listening to the zither.
6. On all other occasions, Kim does not enjoy listening to all other instruments.
7. Everyone else enjoys listening to all other instruments on this occasion.

8. On all other occasions, everyone else has enjoyed listening to all other instruments.

Item 11

1. Alice displayed sculptures made from clay in her home on this occasion.
2. Everyone else displayed sculptures made from clay on this occasion.
3. Alice displayed sculptures made from all other kinds of materials on this occasion.
4. On all other occasions, Alice has displayed sculptures made from clay.
5. On all other occasions, nobody else has displayed sculptures made from clay.
6. On all other occasions, Alice has displayed sculptures made from all other kinds of materials.
7. Everybody else displayed sculptures made from all other kinds of materials on this occasion.
8. On all other occasions, nobody else has displayed sculptures made from all other kinds of materials.

Item 12

1. Fred sang "Golden Slumbers" (a children's song) on this occasion.
2. Everyone else sang "Golden Slumbers" on this occasion.
3. Fred sang all other children's songs on this occasion.
4. On all other occasions, Fred has sung "Golden Slumbers".
5. On all other occasions, everyone else has sung "Golden Slumbers".
6. On all other occasions, Fred has never sung all other children's songs.
7. Everyone else sang all other children's songs on this occasion.
8. On al other occasions, nobody else has sung all other children' songs.

Item 13

1. Eric did not like to dance the tango on this occasion.
2. Nobody else liked to dance the tango on this occasion.
3. Eric did not like to dance all other dances on this occasion.
4. On all other occasions, Eric has never liked to dance the tango.
5. On all other occasions, everyone else has liked to dance the tango.
6. On all other occasions, Eric has liked to dance all other dances.
7. Everyone else liked to dance all other dances on this occasion.
8. On all other occasions, everyone else has liked to dance all other dances.

Item 14

1. Susan did not bow to the statue of fire on this occasion.
2. Nobody else bowed to the statue of fire on this occasion.

3. Susan did not bow to any other statute on this occasion.
4. On all other occasions, Susan has never bowed to the statue of fire.
5. On all other occasions, nobody has ever bowed to the status of fire.
6. On all other occasions, Susan has bowed to all other statues.
7. Everyone else bowed to all other statues on this occasion.
8. On all other occasions, everyone else has always bowed to all other statues.

Item 15

1. George likes collard greens (a leafy vegetable) on this occasion.
2. Everyone else likes collard greens on this occasion.
3. George likes all other kinds of leafy vegetables on this occasion.
4. On all other occasions, George has always liked collard greens.
5. On all other occasions, nobody else has ever liked collard greens.
6. On all other occasions, George has never liked all other kinds of leafy vegetables.
7. Everyone else likes all other kinds of leafy vegetables on this occasion.
8. On all other occasions, nobody else has ever liked all other kinds of leafy vegetable.

Appendix D: materials used for Experiment 4

<u>Kim had a traffic accident last night.</u>
1. Mechanism (person): Kim is near-sighted and tends not to wear her glasses while driving.
2. Covariation (person): Kim is much more likely to have traffic accidents than other people are.
3. Mechanism (occasion): There was a very severe storm and the roads were very slick last night.
4. Covariation (occasion): Traffic accidents were much more likely last night than on other nights.

<u>Joanne was really nervous when she was taking the exam last week.</u>
1. Mechanism (person): Joanne does not know how to prepare for exams.
2. Covariation (person): Joanne is much more likely than an average person to be nervous when taking exams.
3. Mechanism (occasion): It was finals week last week.
4. Covariation (occasion): Last week people were very much more likely to be nervous than they were other weeks.

<u>Mark said the "Seven Glories" prayer at sundown.</u>
1. Mechanism (person): Mark is a member of the clergy.

2. Covariation (person): Mark is very much more likely than other people to say the prayer at sundown.
3. Mechanism (occasion): The rules of Mark's religion suggests that the prayer typically be said at sundown.
4. Covariation (occasion): The prayer is much more likely to be said at sundown than at other times.

Dave got sick to his stomach this morning after eating chicken last night at a local restaurant.

1. Mechanism (stimulus): The chef at the restaurant always undercooks chicken.
2. Covariation (stimulus): People are much more likely to get sick after eating chicken at the restaurant than after eating other foods.
3. Mechanism (person): Dave's stomach lining is easily irritated.
4. Covariation (person): Dave is much more likely than other people to get sick to his stomach.

Yesterday Al went to the Dragons game at the Dragons stadium.

1. Mechanism (stimulus): The Dragons are a really good team.
2. Covariation (stimulus): People are much more likely to go to Dragons games than to go to other games.
3. Mechanism (person): Al sells hot dogs at the Dragons stadium.
4. Covariation (person): Al is much more likely than other people to go to the Dragons stadium.

Sharon didn't enjoy last night's football game.

1. Mechanism (stimulus): The game was pretty boring.
2. Covariation (stimulus): People are somewhat less likely to enjoy this football game than to enjoy other games.
3. Mechanism (person): Sharon thinks that football is too violent.
4. Covariation (person): Sharon is much less likely to enjoy football games than are other people.

Mary did not enjoy dancing with Fred at the annual office party this year.

1. Mechanism (stimulus): Fred is a poor dancer.
2. Covariation (stimulus): People are much less likely to enjoy dancing with Fred than with other people.
3. Mechanism (occasion): The management invites a very bad band to the annual office party.
4. Covariation (occasion): People are much less likely to enjoy dancing with Fred at the annual office party than at other parties.

Tom bought a coat at Briarwood mall last week.

1. Mechanism (stimulus): It was the kind of coat worn by a famous rock star.
2. Covariation (stimulus): People are somewhat more likely to buy these coats than they are to buy other kinds of coat.
3. Mechanism (occasion): Last week was the coldest week of the year.
4. Covariation (occasion): People were somewhat more likely to buy these coats last week than they were to buy them other weeks.

Ellen didn't drink the French wine at dinner.
1. Mechanism (stimulus): The French wine was cheap and not too good.
2. Covariation (stimulus): At the dinner, people were a bit less likely to drink the French wine than other kinds of wine.
3. Mechanism (occasion): It was a business dinner where people were negotiating.
4. Covariation (occasion): People were barely less likely to drink the French wine at dinner than they were to drink it on other occasions.

Appendix E: test of a different model for Experiment 4

It may be questioned whether the low R-squares were due to the inappropriateness of our constraint that supporting and competing information were only allowed to differ in sign. In other words, perhaps supporting and competing information have different magnitudes of effect. We tested the fit of two regressions with separate variables for supporting and competing versions of the same kind of information (i.e., predic tors: mechanism-for, mechanism-against, covariation-for, and covariation-against). Inspection of the confidence intervals for the beta weights of these models shows that we cannot reject the hypothesis that the for and against variables of the same kind of information had the same magnitude of effect (i.e., only differed in sign; regression 1 90% intervals: mechanism-for, −.9 to −.46, mechanism-against, .75 to 1.2, covariation-for, −.24 to −.69, covariation-against, .04 to .49; regression 2 90% intervals: mechanism-for, −.62 to −1.1, mechanism-against, .41 to .89, covariation-for, −.03 to −.5, covariation-against, .1 to .57). Additionally, sums of squares for the four predictor models were almost identical to those of the direct competition models (regression 1: direct-SS, 40.1 vs. 4-predictor-SS, 44.4; regression 2: direct-SS, 34.3 vs. 4-predictor-SS, 35.5).

References

Abelson, R.P., & Lalljee, M. (1988). Knowledge structures and causal explanation. In D. Hilton (Ed.), *Contemporary science and natural explanation: commonsense conceptions of causality*. Brighton: Harvester Press.

Ahn, W., Brewer, W.F., & Mooney, R.J. (1992). Schema acquisition from a single example. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *18*, 391–412.

Barsalou, L.W. (1983). Ad hoc categories. *Memory and Cognition*, *11*, 211–227.
Bullock, M., Gelman, R., & Baillargeon, R. (1982). The development of causal reasoning. In W.J. Friedman (Ed.), *The developmental psychology of time* (pp. 209–254). San Diego, CA: Academic Press.
Chapman, L.J., & Chapman, J.P. (1967). Genesis of popular but erroneous diagnostic observations. *Journal of Abnormal Psychology*, *72*, 193–204.
Chapman, L.J., & Chapman, J.P. (1969). Illusory correlation as an obstacle to the use of valid psychodiagnostic signs. *Journal of Abnormal Psychology*, *74*, 272–280.
Cheng, P.W. (1993). Separating causal laws from causal facts: pressing the limits of statistical relevance. In D.L. Medin (Ed.), *The psychology of learning and motivation* (Vol. 30, pp. 215–264). New York: Academic Press.
Cheng, P.W., & Novick, L.R. (1990a). A probabilistic contrast model of causal induction. *Journal of Personality and Social Psychology*, *58*, 545–567.
Cheng, P.W., & Novick, L.R. (1990b). Where is the bias in causal attribution? In K. Gilhooly, M. Keane, R. Logie, & G. Erdos (Eds.), *Lines of thought: reflections on the psychology of thinking* (pp. 181–197). Chichester: Wiley.
Cheng, P.W., & Novick, L.R. (1991). Causes versus enabling conditions. *Cognition*, *40*, 83–120.
Cheng, P.W., & Novick, L.R. (1992). Covariation in natural causal induction. *Psychological Review*, *99*, 365–382.
Einhorn, H.J., & Hogarth, R.M. (1986). Judging probable cause. *Psychological Bulletin*, *99*, 3–19.
Gelman, S.A., & Medin, D.L. (1993). What's so essential about essentialism? A different perspective on the interaction of perception, language, and conceptual knowledge. *Cognitive Development*, *8*, 157–167.
Gelman, S.A., & Wellman, H.M. (1991). Insides and essence: early understandings of the non-obvious. *Cognition*, *38*, 213–244.
Gopnik, A., & Wellman, H.M. (1994). The theory theory. In L. Hirschfeld & S. Gelman (Eds.), *Domain specificity in cognition and culture*. New York: Cambridge University Press.
Harré, R., & Madden, E.H. (1975). *Causal powers: a theory of natural necessity*, Totowa, NJ: Rowman & Littlefield.
Hart, H.L.A., & Honoré, A.M. (1959). *Causation in the law*. Oxford: Clarendon Press.
Hewstone, M.R.C. (1989). *Causal attribution*, Cambridge, MA: Basil Blackwell.
Hewstone, M.R.C., & Jaspars, J.M.F. (1987). Covariation and causal attribution: a logical model of the intuitive analysis of variance. *Journal of Personality and Social Psychology*, *53*, 663–672.
Hilton, D.J. (1988). *Contemporary science and natural explanation: commonsense conceptions of causality*. New York: New York University Press.
Hilton, D.J. (1990). Conversational processes and causal explanation. *Psychological Bulletin*, *107*, 65–81.
Hilton, D.J., & Slugoski, B.R. (1986). Knowledge-based causal attribution: the abnormal conditions focus model. *Psychological Review*, *93*, 75–88.
Hume, D. (1978). *A treatise of human nature*. Oxford: Oxford University Press. (Original work published 1739.)
Jaspars, H.M.F., Hewston, M., & Fincham, F.D. (1983). Attribution theory and research: the state of the art. In J.M.F. Jaspars, F.D. Fincham, & M. Hewstone (Eds.), *Attribution theory and research: conceptual, developmental and social dimensions*. London: Academic Press.
Kassin, S.M. (1979). Consensus information, prediction and causal attribution: a review of the literature and issues. *Journal of Personality and Social Psychology*, *37*, 1966–1981.
Kelley, H.H. (1967). Attribution theory in social psychology. In D. Levine (Ed.), *Nebraska symposium on motivation* (Vol. 15, pp. 192–238). Lincoln, NB: University of Nebraska Press.
Kelley, H.H. (1973). The processes of causal attribution. *American Psychologist*, *28*, 107–128.

Kitcher, P., & Salmon, S.C. (1989). *Minnesota studies in the philosophy of science: Vol. XIII. Scientific explanation*. Minneapolis: University of Minnesota Press.

Koslowski, B., Spilton, D., & Snipper, A. (1981). Children's beliefs about instances of mechanical and electrical causation. *Journal of Applied Developmental Psychology, 2*, 189–210.

Lalljee, M. & Abelson, R. (1983). The organization of explanations. In M. Hewstone (Ed.), *Attribution theory: social and functional extensions*. Oxford: Basil Blackwell.

Lalljee, M., Lamb, R. Furnham, A., & Jaspars, J.M.F. (1984). Explanations and information search: inductive and hypotheses-testing approaches to arriving at an explanation. *British Journal of Social Psychology, 23*, 201–212.

Leddo, J. & Abelson, R.P. (1986). The nature of explanations. In J.A. Galambos, R.P. Abelson and J.B. Black (Eds.), *Knowledge structures*. Hillsdale, NJ: Erlbaum.

Lipe, M.G. (1985). *Attribution theory: a proposed model*, unpublished doctoral dissertation, University of Chicago.

Lipe, M.G. (1991). Counterfactual reasoning as a framework for attribution theories. *Psychological Bulletin, 109*, 456–471.

Mackie, J.L. (1975). Causes and conditionals. In E. Sosa (Ed.), *Causation and conditionals*. Oxford: Oxford University Press.

Major, B. (1980). Information acquisition and attribution processes. *Journal of Personality and Social Psychology, 39*, 1010–1023.

Malt, B.C. (1990). Features and beliefs in the mental representation of categories. *Journal of Memory and Language, 29*, 289–315.

McArthur, L.A. (1972). The how and what of why: some determinants and consequences of causal attribution. *Journal of Personality and Social Psychology, 22*, 171–193.

Medin, D.L. (1989). Concepts and conceptual structure. *American Psychologist, 44*, 1469–1481.

Medin, D.L., & Ortony, A. (1988). Psychological essentialism. In S. Vosniadou & A. Ortony (Eds.), *Similarity and analogical reasoning* (pp. 179–196). New York: Cambridge University Press.

Michotte, A.E. (1963). *The perception of causality*. New York: Basic Books. (Original work published 1946.)

Mill, J.S. (1973). A system of logic ratiocinative and inductive. In J.M. Robson (Ed.), *Collected works of John Stuart Mill* (Vols. 7, 8).Toronto: University of Toronto Press. (Original work published 1843.)

Murphy, G.L., & Medin, D.L. (1985). The role of theories in conceptual coherence. *Psychological Review, 92*, 289–316.

Pappineau, D. (1989). Pure, mixed, and spurious probabilities and their significance for a reductionist theory of causation. In P. Kitcher & W.C. Salmon (Eds.), *Minnesota studies in the philosophy of science: Vol. XIII. Scientific explanation* (pp. 307–348). Minneapolis: University of Minnesota Press.

Pazzani, M.J. (1991). Influence of prior knowledge concept acquisition: experimental and computational results. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 17*, 416–432.

Salmon, W.C. (1970). Statistical explanation. In W. Salmon, R. Jeffrey, & J. Greeno (Eds.), *Statistical explanation and statistical relevance*. Pittsburgh: University of Pittsburgh Press.

Salmon, W.C. (1984). *Scientific explanation and the causal structure of the world*. Princeton, NJ: Princeton University Press.

Shaver, K.G. (1981). Back to basics: on the role of theory in the attribution of causality. In H. Harvey, W.J. Ickes and R.F. Kidd (Eds.). *New directions in attribution research* (Vol. 3). Hillsdale, NJ: Erlbaum.

Shultz, T.R. (1982). Rules of causal attribution. *Monographs of the Society for Research in Child Development, 47* (Serial No. 194).

Shultz, T.R. (1986). Selection of causal rules. *Child Development, 57*, 143–152.

Shultz, T.R. Fisher, G.W., Pratt, C.C., & Rulf, S. (1986). Selection of causal rules. *Child Development*, *57*, 143–152.

Spirtes, P., Glymour, C., & Scheines, R. (1993). *Causation, prediction, and search*. Springer-Verlag, New York.

Suppes, P. (1970). *A probabilistic theory of causality*. Amsterdam: North-Holland.

White, P.A. (1990). Ideas about causation in philosophy and psychology. *Psychological Bulletin*, *108*, 3–18.