ATTITUDES AND SOCIAL COGNITION

A Probabilistic Contrast Model of Causal Induction

Patricia W. Cheng University of California, Los Angeles Laura R. Novick Vanderbilt University

Deviations from the predictions of covariational models of causal attribution have often been reported in the literature. These include a bias against using consensus information, a bias toward attributing effects to a person, and a tendency to make a variety of unpredicted conjunctive attributions. It is contended that these deviations, rather than representing irrational biases, could be due to (a) unspecified information over which causal inferences are computed and (b) the questionable normativeness of the models against which these deviations have been measured. A probabilistic extension of Kelley's analysis-of-variance analogy is proposed. An experiment was performed to assess the above biases and evaluate the proposed model against competing ones. The results indicate that the inference process is unbiased.

As early as the age of 3, children form their own hypotheses about the causes of events happening in their world (e.g., Bullock, Gelman, & Baillargeon, 1982). The child's preoccupation with understanding the causes of events continues into adulthood, presumably because this knowledge is required for successful interaction with the environment. Assuming that the goals of causal explanation are the prediction and potential control of future events, it would seem important for people to induce causes in a rational (i.e., normative) manner.

Work on causal attribution, in fact, began with the proposal that causal induction is normative. On the basis of earlier work by Heider (1958), Kelley (1967) proposed that people are intuitive scientists. According to Kelley's influential model, causal inferences are based on the principle that he termed *covariation*. The principle can be traced back to Mill's (1843/1973) "joint method of agreement and difference," an inductive method that Mill proposed as normative. More specifically, Kelley proposed that causal inferences are based on a statistical interpretation of the covariation principle as instantiated in the

analysis of variance (ANOVA). Analysis of variance is, of course, generally regarded as a normative inductive procedure.

Subsequent research on causal attribution, however, has identified many deviations from the predictions of covariational models. These deviations include (a) a bias against using consensus information, (b) a bias toward attributing effects to a person, and (c) a tendency to make a variety of other unpredicted attributions, particularly conjunctive ones (e.g., see Borgida & Brekke, 1981; Hilton & Slugoski, 1986; Jaspars, Hewstone, & Fincham, 1983; Nisbett & Ross, 1980).

Is the causal attribution process inherently biased? We argue that to answer this question it is important to distinguish between the data on which the causal inference process operates and the process of inference computation itself. Just as false conclusions in deductive reasoning can be reached by the use of valid deductive rules when the premises on which the rules operate are false (Henle, 1962), so observed biases in inductive reasoning may be due to the nature of the input (i.e., the set of information on which inference is computed and the pattern of that information) rather than to biases in the process of inference computation itself. Previous work on causal attribution typically has not made this distinction or has not accurately identified the data on which induction operates. Consequently, in most cases it is not possible to determine whether observed biases occur in the inference process per se or in the data on which the inference process operates.

Even when the pattern of information on which inference is computed is identified, a model must be normative if deviations from it are to be considered "biases." We suggest that some deviations from previous models may, in fact, be rational inferences, and we propose a normative model against which to reassess biases—our *probabilistic contrast model.*¹ It is a covariational

The research reported in this article was supported by Grant BNS87-10305 from the National Science Foundation to Patricia W. Cheng and by a National Institute of Mental Health Individual Postdoctoral Fellowship to Laura R. Novick. The results reported here were presented in part at the International Conference on Thinking, Aberdeen, Scotland, August 1988.

We thank Paul Abramson, Frederick Försterling, Rochel Gelman, Denis Hilton, Keith Holyoak, Edward Jones, Harold Kelley, Yunn-wen Lien, Richard Nisbett, Craig Smith, Eliot Smith, Bernard Weiner, and two anonymous reviewers for their helpful comments on an earlier draft. We also thank Todd Gross and especially Livian Perez for their technical assistance.

Correspondence concerning this article should be addressed to Patricia W. Cheng, Department of Psychology, 1285 Franz Hall, University of California, Los Angeles, California 90024-1563.

¹ In an earlier article, we (Cheng & Novick, 1990) referred to this model as the "qualitative contrast model."

Journal of Personality and Social Psychology, 1990, Vol. 58, No. 4, 545-567 Copyright 1990 by the American Psychological Association, Inc. 0022-3514/90/\$00.75

model based on estimated differences in the probabilities of the effect conditional on the presence versus the absence of potential causal factors. Our model can be viewed as an extension of Kelley's (1967, 1973) analogy between causal induction and analysis of variance and of Jenkins and Ward's (1965) contingency rule. (We discuss the relationship between our model and related models from the cognitive psychology literature in Cheng & Novick, 1989a.) Although Kelley's ANOVA model has inspired an enormous amount of research, the analogy between the model and an actual analysis of variance has been left relatively vague. This vagueness has allowed interpretations of the model that were, in fact, inconsistent with the analogy (e.g., Mc-Arthur, 1972; Orvis, Cunningham, & Kelley, 1975; Pruitt & Insko, 1980), as we will discuss. Our model formulates the analogy without assuming the implausible mental computations underlying the analysis of variance. In the final section of this article, we report an experiment that reassesses the above-mentioned biases and tests our model against competing ones (Försterling, 1989; Hilton & Slugoski, 1986; Jaspars et al., 1983; Orvis et al., 1975) in light of types of additional information that typically have been ignored in experiments on causal attribution.

Normative Covariation-Based Models

Before we present our explanation of the biases, we will briefly review two models of causal attribution that may be regarded as both descriptive and normative. They represent different interpretations of Kelley's (1967, p. 194) covariation principle, which states: "The effect is attributed to that condition which is present when the effect is present and which is absent when the effect is absent." The models based on these interpretations are the ones against which the biases have been measured.

Kelley's Analysis of Variance Model

Kelley (1967, 1973) interpreted his covariation principle statistically in his ANOVA model. He proposed the dimensions of persons (P), stimuli (S), and time/modalities (T) as independent variables in the model, as illustrated in his Persons × Entities \times Time/Modalities cube. To measure covariation along these respective dimensions, Kelley proposed three information variables, which he referred to as consensus, distinctiveness, and consistency. According to this model (Kelley, 1967; Kelley & Michela, 1980), one infers the cause of a given P's response to a certain S on a particular occasion T based on one's perception of the degree of (a) consensus between P's response to S and other people's responses to S (on occasion T), (b) distinctiveness of P's response to S from P's responses to other stimuli (on occasion T), and (c) consistency of P's response to S on occasion T with P's responses to S on other occasions. Note that consensus and consistency are inversely proportional to covariation (e.g., high consensus indicates low covariation between the particular person and the effect), whereas distinctiveness is directly proportional to covariation.

In addition to proposing this general model, Kelley made specific attributional predictions for three particular configurations of consensus, distinctiveness, and consistency information: (a) Low consensus, low distinctiveness, and high consistency (LLH) should lead to a person attribution; (b) high consensus, high distinctiveness, and high consistency (HHH) should lead to a stimulus attribution; and (c) low consensus, high distinctiveness, and low consistency (LHL) should lead to a circumstance attribution. (We will use the term *configuration* to denote the pattern of information on the three variables of consensus, distinctiveness, and consistency.) In terms of covariation, the LLH configuration, for example, indicates high covariation along the person dimension and low covariation along the other two dimensions, thus leading to a "main-effect" attribution of person in Kelley's (1973) terms.

An extension of the ANOVA model was proposed by Orvis et al. (1975), according to which the three configurations described above, considered more fundamental than other configurations, are used as templates on the basis of which predictions for other configurations are made. Consider, for example, a situation involving high consensus, low distinctiveness, and low consistency (HLL). The high consensus in this configuration matches only the HHH template, thus suggesting a stimulus attribution. The low distinctiveness matches only the LLH template, suggesting a person attribution. Finally, the low consistency matches only the LHL template, suggesting a circumstance attribution. Because each of these three attributions is suggested with equal frequency, the template model predicts that the cause will be attributed to a combination of the three attributions, with the combination being in the form of an "and/or" relation (i.e., person and/or stimulus and/or circumstances; see Orvis et al., 1975, p. 607). As Jaspars et al. (1983) and Försterling (1989) have noted, because attributions for some configurations are based on a pattern-matching process, the template model is not consistent with the ANOVA analogy.

Jaspars, Hewstone, and Fincham's Inductive Logic Model

Jaspars et al. (1983; also see Hewstone & Jaspars, 1987) proposed a deterministic interpretation of Kelley's (1967) covariation principle in what they termed their "inductive logic model" for causal attribution. Like Kelley, Jaspars et al. applied their model to configurations of consensus, distinctiveness, and consistency information. In this model, each piece of information given in an attribution problem is coded in terms of whether it involves (a) the person in question or other people, (b) the stimulus in question or other stimuli, and (c) the occasion in question or other occasions. In addition, the presence or absence of the effect (i.e., the behavior in question) is noted. According to this model, people note for each possible causal factor (or conjunction of factors) whether it is present when the effect is present and absent when the effect is absent, and identify a factor (or conjunction of factors) that is both necessary and sufficient for the occurrence of the effect as the cause. This model instantiates Mill's "joint method" for the three dimensions of persons, stimuli, and occasions based on configurational information.

Hypothesized Biases

Several findings pose problems for the two covariation-based models discussed above. We briefly review the person and consensus biases that have been well-documented in the literature, as well as various other deviations from these models.

Person and Consensus Biases

Although Kelley's (1967, 1973) predictions have received strong support from several major studies (McArthur, 1972, 1976; Orvis et al., 1975), these studies have also revealed some large deviations from the predictions. In this section, we consider the bias against using consensus information relative to information on the other two variables and the related bias toward attributing effects to a person.

In an experiment that orthogonally varied two levels (high versus low) of consensus, distinctiveness, and consistency, Mc-Arthur (1972) found that consistency information accounted for 41% of the variance in circumstance attributions and distinctiveness accounted for 12% of the variance in stimulus attributions, but consensus information accounted for only 6% of the variance in person attributions. This same underuse of consensus information was observed in the prediction of total variance in causal attributions: Consistency accounted for 20% of the variance, distinctiveness accounted for 10% of the variance, and consensus accounted for 3% of the variance. This bias against using consensus information has been found in numerous other studies (for reviews, see Alloy & Tabachnik, 1984; Borgida & Brekke, 1981; and Kassin, 1979), leading some investigators to argue that it is a general phenomenon (e.g., Borgida, 1978; Nisbett & Borgida, 1975). To complicate the picture, however, other work has found that the consensus bias does not occur under all conditions (Hansen & Donoghue, 1977; Hewstone & Jaspars, 1983; Kulik & Taylor, 1980; Pruitt & Insko, 1980; Ruble & Feldman, 1976; Weiner et al., 1971; Wells & Harvey, 1977; Zuckerman, 1978).

The bias toward attributing effects to a person rather than to a stimulus or to circumstances has likewise appeared in many studies (e.g., McArthur, 1972; Orvis et al., 1975; Pruitt & Insko, 1980). For example, Jaspars et al. (1983) noted that 82% of Mc-Arthur's (1972) subjects attributed an effect to the person when the presence of the person was both necessary and sufficient to produce the effect, compared with 62% stimulus attributions and 33% circumstance attributions when those factors were necessary and sufficient. This bias toward person attributions is consistent with what Ross (1977) has called the "fundamental attribution error." Like the consensus bias, however, the person bias does not appear in all situations. For example, Hilton and Jaspars (1987) found a bias toward attributing effects to the stimulus instead. Yet another pattern was found in Försterling's (1989, Table 5) data, indicating that person, stimulus, and time attributions were equally likely to be made, with mean ratings of 8.2, 8.0, and 8.5 (on a 10-point scale), respectively, when each of those factors was both necessary and sufficient for the effect to occur.

The above two biases, to the extent that they occur, indicate apparent deviations from the predictions of any purely covariation-based model such as Kelley's ANOVA model or Jaspars et al.'s inductive logic model, because the amounts of covariation along the three dimensions were presumably counterbalanced in the set of stimuli.

Deviations From Predictions for Particular Configurations

Kelley (1967, 1973) made predictions for only three of the eight possible information configurations. In contrast, Jaspars et al. (1983; also see Hewstone & Jaspars, 1987) made predictions for all eight configurations. Their prediction for one of these additional configurations is a particularly clear standard against which biases may be judged: For the high consensus, low distinctiveness, high consistency (HLH) configuration, the inductive logic model predicts that no causal attribution is possible, because no factor covaries with the effect. This prediction, which follows directly from the application of the covariation principle to the configuration, is consistent with the spirit of Kelley's ANOVA model, even though Kelley did not explicitly make the prediction.

Contrary to this prediction, the HLH configuration has been a locus of both of the biases described above. For this configuration, subjects in several studies showed a clear preference for attributing the effect to the person, despite the lack of covariation of the person with the effect. For example, 45% of the subjects in McArthur's (1972) study, 43% of the subjects in Study 2 of Orvis et al. (1975), and 23% of the subjects in Jaspars's (1983) study made a person attribution. In addition, subjects often attributed the effect to the stimulus (19% of the subjects in Study 2 of Orvis et al.; 20% of Jaspars's subjects; and 31% of Hilton & Jaspars's, 1987, subjects) or to the combination (i.e., conjunction) of the person and the stimulus (35% of McArthur's subjects; 27% of Jaspars's subjects; and 24% of Hilton & Jaspars's subjects). Similarly, Orvis et al. found many attributions to "the person and/or the stimulus" (56% and 31% of the subjects in Studies 1 and 2, respectively).

Deviations from Jaspars et al.'s predictions for other configurations have also been found. Consider the high consensus, low distinctiveness, low consistency (HLL) configuration, for which the inductive logic model predicts an attribution to the time dimension.² In the studies of Hilton and Slugoski (1986) and Försterling (1989), subjects were asked to rate each main effect and interaction attribution for how complete an explanation it provided for the target event. The predicted occasion attribution received a mean rating of 5.6 on a 7-point scale in Hilton and Slugoski's data and a mean rating of 8.5 on a 10-point scale in Försterling's data. In Hilton and Jaspars's (1987) study, the circumstance attribution was chosen by 19% of the subjects. In all three sets of data, however, conjunctive attributions also were perceived as important causal factors. In Försterling's data, the stimulus-and-occasion attribution received a mean rating of 7.7. In Hilton and Slugoski's data, that attribution received a mean rating of 5.1. In addition, the person-and-occasion attribution received a mean rating of 5.7. The corresponding per-

² The circumstance attribution is ambiguous because it could be interpreted by subjects either as a simple effect of occasion or as some interactive combination of the person, stimulus, and occasion. Because of this ambiguity, we only report the results of those studies for which the response form included either (a) interactions of the circumstances with persons and stimuli in addition to the simple circumstance attribution or (b) the unambiguous "occasion" attribution (and interactions of the other variables with that variable).

son-and-circumstance attribution was chosen by 20% of Hilton and Jaspars's subjects. Conjunctive attributions have also been found for the LLH and HHH configurations (e.g., Hilton & Jaspars, 1987; Hilton & Slugoski, 1986; Jaspars, 1983; Mc-Arthur, 1972), for which main effects of person and stimulus, respectively, have been predicted.

An Informational Explanation of the Biases

Why are there such biases, and why do they capriciously appear under some conditions but not others? Many explanations have been proposed to answer these questions. A common one is that prior expectations sometimes override presumably more objective data-based processing (Alloy & Tabachnik, 1984; Kassin, 1979; Orvis et al., 1975; Tversky & Kahneman, 1982). In discussing the deviations from Jaspars et al.'s (1983) model for the HLH configuration, Hilton and Slugoski (1986) took the more radical view that the biases do not stem from amendments to the covariational criterion of causality, but rather from subjects' use of a different criterion altogether-the criterion of an "abnormal condition." Recently, we (Cheng & Novick, 1990) proposed an informational explanation for the person and consensus biases that retains the position that causal induction is based on an assessment of covariation. We now extend this explanation to account for the additional biases described above.

Recall that earlier we made a distinction between the inference process itself and the information on which the inference process operates (i.e., the input to that process). Most experiments testing covariation-based models have specified the input in terms of the variables of consensus, distinctiveness, and consistency. As noted by several investigators (Cheng & Novick, 1990; Försterling, 1989; Hilton, 1988, 1990; Jaspars et al., 1983; Pruitt & Insko, 1980), information on these three variables, apparently often assumed to represent all the data relevant to making causal attributions, actually covers only a subset of the potentially relevant information. For example, consensus information indicates the amount of agreement between the person in question and other people in their responses to a particular stimulus on a particular occasion, rather than the amount of response agreement between the person and other people with respect to all the stimuli on all the occasions. Thus, consensus information covers only Region 1 in Figure 1. Similarly, distinctiveness and consistency information respectively cover Regions 2 and 3. The figure demarcates eight regions of information in Kelley's (1967, 1973) Persons × Stimuli × Occasions cube. Three vertical slices of the cube are represented separately, with each slice representing a different occasion. Region 0 denotes the target event to be explained.

In experiments specifying only configurational information, it is typically not known what assumptions subjects might have made spontaneously regarding the occurrence of the effect in the nonconfigurational parts of the cube (Regions 4–7 in Figure 1). We propose (also see Cheng & Novick, 1990) that because causal attribution is a function of the data on which the inference rules operate, as well as of the rules themselves, the apparent biases found in previous experiments—rather than being due to the inferential process—may reflect the assumptions made by subjects regarding the pattern of information for the unspecified cells of the cube. The cognitive literature indicates that people often go beyond the information given in arriving at an analysis of a situation (e.g., Bruner, 1957; Harris, 1977; Johnson, Bransford, & Solomon, 1973). In the causal attribution literature, this conclusion is supported by existing data indicating that subjects do make use of nonconfigurational information. Hilton and Slugoski (1986) reported that people's causal attributions were influenced by their implicit knowledge of norms (i.e., "presuppositions about what a *class* of persons generally does to a *class* of stimuli"; Hilton, Smith, & Alicke, 1988, p. 531). Such knowledge fills the rest of the cube shown in Figure 1 homogeneously with the same pattern of information. Similarly, Pruitt and Insko (1980) reported that "comparisonobject consensus" (i.e., how other people respond to other stimuli on this occasion) influenced causal attributions. Comparison-object consensus fills Region 6.

Given that subjects seem to make use of explicit or implicit information over nonconfigurational parts of the cube in deriving causal attributions, it is possible that when one takes information over the entire cube into account, there may be no systematic *inferential* biases at all. The experiment we report tests this hypothesis.

Previous Approaches to the Incompleteness of Information

Investigators who have noted the incompleteness of configurational information have taken diverse approaches regarding the issue. According to Jaspars et al. (1983, p. 10), "[This limited information] represents an incomplete, fractional replication design in which the independent variables are not orthogonal. Even an explicit analysis of variance with the aid of a computer programme would not be easy." They concluded that subjects must be doing something other than an implicit ANOVA and proposed their inductive logic model, which derives attributions based entirely on configurational information.

Pruitt and Insko (1980) indicated that consideration should be given to the consistency of (a) the other-people-to-target-object relation (Region 4 in Figure 1), (b) the target-person-tocomparison-object relation (Region 5), and (c) the other-peopleto-comparison-object relation (Region 7), besides comparisonobject consensus (Region 6). However, they did not study all of these regions of the cube, choosing to add only comparisonobject consensus to the conventional configuration in their study. The information presented to their subjects was, therefore, still incomplete for an analysis of variance. Their findings were analyzed in terms of the effects of each of the four information variables rather than in terms of an analysis of variance using the three dimensions of persons, stimuli, and occasions as independent variables. Thus, their model, although presented as an extension of Kelley's ANOVA model, does not make use of the analogy to analysis of variance (also see Försterling, 1989). (For the same reason, McArthur's, 1972, predictions are not based on the ANOVA analogy.)

Hilton (1988, 1990) and Försterling (1989) have taken very similar approaches regarding this issue. Both noted that the completion of information would enable the computation of an analysis of variance. Hilton recast his previous findings with Slugoski (Hilton & Slugoski, 1986) in terms of an analysis of variance, illustrating the effect of norms in terms of a two-way



Figure 1. The eight regions of information in Kelley's (1967) Persons \times Stimuli \times Occasions cube. (The shaded regions indicate configurational information.)

analysis of variance (assuming high consistency for all information). Similarly, Försterling presented an interpretation of the ANOVA analogy applied to information over the entire cube. His interpretation, as we will discuss, is equivalent to the deterministic principle underlying Jaspars et al.'s (1983) inductive logic model. Försterling differs from Jaspars et al. only in that he applied the principle to the entire cube rather than to configurational information alone. Försterling's experiment, using materials involving deterministic covariation for a $2 \times 2 \times 2$ cube, generally supported his model.

The Probabilistic Contrast Model of Causal Induction

As mentioned earlier, even when the pattern of information on which inference is computed is identified, whether deviations from a particular model are biases depends on whether the model is normative. We propose a model against which to reassess biases—our probabilistic contrast model. Our work builds on the ideas of Kelley (1967, 1973), Hilton (1988, 1990), and Försterling (1989). Extending the work of Hilton and Slugoski (1986), we address all of the biases discussed earlier, as well as the apparent capriciousness of these biases, within the framework of our model and our distinction between inference rules and data.

One limitation of previous models is that the analogy between causal induction and an analysis of variance has been left vague. For example, theorists (e.g., Försterling, 1989; Kelley, 1973) have proposed simply that people perform a "naive" version of the statistical ANOVA. Kelley (1973, p. 109) suggested that the "naive ANOVA is a poor replica of the scientific one," but we are not aware of any specification of how the naive AN-OVA compares with the statistical procedure. Our model is such a specification. It is an explicit model of causal induction that abstracts from the ANOVA analogy its key elements, avoiding the assumption of the complex quantitative computations underlying the statistical procedure (e.g., sums of squares and *F* ratios).

Our model of causal induction is a probabilistic analogue of statistical contrasts. It is termed a *contrast* model because it refers to contrasts between a particular value on a dimension and other values on that dimension, or to contrasts involving particular combinations of values as opposed to other combinations of values. These are intended to reflect, respectively, responses such as "There is something special about the person" and "There is something special about the combination of this person and this stimulus." Such responses correspond to specific contrasts rather than to general main effects or interactions in the analysis of variance. We do not believe that people mentally perform the complex quantitative computations underlying statistical contrasts. Therefore, our model only requires that people be able to estimate and compare proportions, a task that has been found to be performed reasonably well by naive subjects (Alloy & Abramson, 1979; Robinson, 1964; Shuford, 1961). Our model accounts for causal induction based on qualitative as well as quantitative data, as we will explain when we make predictions for the stimuli in our experiment.

Contrasts are assumed to be computed for attended dimensions that are present in the event to be explained. For maineffect contrasts, we define a cause to be a factor for which P[i], the proportion of cases for which the effect occurs when factor i is present, is greater (by some criterion) than $P[\sim i]$, the proportion of cases for which the effect occurs when factor i is absent. (These proportions are estimates of the probabilities of the effect occurring conditional on the presence of i and the absence of i, respectively. A tilde before a factor label denotes the absence of that factor.) In other words, a cause is a factor, the presence of which (relative to its absence) increases the likelihood of the effect. The criterion is an empirically determined parameter that defines a noticeable difference. The magnitude of the criterion should reflect people's awareness of random sampling fluctuations as well as the role of sample size in people's interpretations of such fluctuations (e.g., see Nisbett, Krantz, Jepson, & Kunda, 1983). We leave investigation of this parameter for future research. In this article we focus our discussion on clear-cut contrasts. We assume that the sample of cases on each dimension is relatively large and that the criterion for differences in proportions is not too small, to rule out situations for which differences are distinguishable only by statistical methods. To illustrate a main-effect contrast in terms of Kelley's Persons × Stimuli × Occasions cube, if the proportion of cells in which the target effect occurs for Person 1 (the top plane in Figure 2) is greater than the corresponding proportion for other people (the rest of the cube), then a target effect involving Person 1 will be attributed to Person 1.

So far, we have only considered causes that consist of a single



Figure 2. Information relevant for specifying an interaction contrast according to the probabilistic contrast model.

factor (e.g., the effect is due to something about a particular person). Sometimes, however, effects are due to a conjunction of factors. For example, someone's allergy to a substance may be attributable to the conjunction (i.e., interaction) of the particular person and the particular substance. Whereas a maineffect contrast in our model requires a noticeable difference between the proportions of cases in which the effect occurs in the presence of a factor and in the absence of it, an *interaction* contrast requires a noticeable differences for levels of an orthogonal factor. In other words, it requires a second-order difference.

Consider Figure 2 again. Let Person i be a particular value along the person dimension (i = 1 in Figure 2) and Stimulus j be a particular value along the stimulus dimension (j = 1) in Figure 2). Further suppose that we wish to determine whether the conjunction of Person i and Stimulus j would be considered a cause for a target event involving Person i, Stimulus j, and Occasion k. A comparison involving four proportions is relevant: For events in which Person i is present (i.e., the top plane in the figure), let P[i,j] be the proportion of occasions (cells along the remaining dimension) on which the effect occurs when Stimulus j is present (the darkly shaded beam in the figure) and let P[i,~j] be the corresponding proportion when Stimulus j is absent (the lightly shaded area); for events in which Person i is absent (i.e., the bottom four planes), let P[~i,j] be the proportion of occasions on which the effect occurs when Stimulus j is present (the striped area) and $P[\sim i, \sim j]$ be the corresponding proportion when Stimulus j is absent (the unshaded area). If the difference between P[i,j] and $P[i, \sim j]$ is noticeably greater than the difference between $P[\sim i, j]$ and $P[\sim i, \sim j]$, or, equivalently, if the difference between P[i,j] and P[~i,j] is noticeably greater

than the difference between $P[i, \sim j]$ and $P[\sim i, \sim j]$, then the conjunction (combination) of Person i and Stimulus j is a cause (e.g., if the proportion of cells in which the target effect occurs is large in the darkly shaded beam but the proportion for the rest of the cube is uniformly small). Interaction contrasts involving the other two combinations of dimensions are defined analogously. More generally, interaction contrasts involving nfactors are defined as nth order differences, where n is any positive integer. (People will no doubt have greater difficulty with interaction contrasts involving greater complexity, and at some maximum level of complexity computation presumably will become impossible. However, because the probabilistic contrast model is a computational model [in Marr's, 1982, sense of the term] that specifies what is computed, rather than a process model that specifies how the computation is carried out, it does not deal with such processing limitations. It seems reasonable to expect that a model of processing limitations should apply across many different types of tasks rather than being specific to inference tasks.)

The causal factors we have considered so far have been facilitatory factors, that is, factors that increase the likelihood of the effect. Analogously, there are also inhibitory factors, which decrease the likelihood of an effect. For example, the presence of a vaccine in one's body decreases the likelihood of one contracting a particular disease (also see Hilton & Slugoski, 1986, and Kelley, 1973, on this distinction). Our probabilistic contrast model distinguishes between inhibitory and facilitatory factors (unlike the standard ANOVA, which squares all differences and therefore loses all information about directionality). In our model, the difference in proportions is positive for facilitatory main-effect contrasts and negative for inhibitory main-effect contrasts. For interaction contrasts, the difference between the differences is positive for facilitatory combinations of factors and negative for inhibitory combinations.

Note that our model clearly allows for multiple alternative causes, and distinguishes them from conjunctive causes (i.e., causes involving multiple necessary factors). Any factor whose presence increases the probability of an effect is a possible cause of the effect. If multiple factors are observed to increase the probability of the effect, then each is a cause of the effect and none is necessary. The multiple alternative causes, of course, can be any combination of main-effect causes and/or conjunctive causes (i.e., interactions). We will provide examples of situations with multiple alternative causes when we describe our stimuli.

To summarize, our model is a probabilistic interpretation of the covariation principle. More specifically, it is an analogue of statistical contrasts that requires only the assessment and comparison of (quantitative or qualitative) proportions. Like its statistical counterpart, however, it is a general procedure that is not committed to particular dimensions (e.g., the three in Kelley's cube) or even to any particular number of dimensions.³

³ Our model assumes that a cause must be perceived as preceding its effect. In other words, covariation should be computed only for factors that are perceived to be temporally prior to the effect. We assume perceived temporal priority rather than manipulation as an augmenting criterion because manipulation, although a standard criterion in many sciences, seems too strong a requirement for everyday induction. Mi-

EXPERIMENT

The two purposes of our experiment are (a) to evaluate our probabilistic contrast model against alternative models of causal attribution (Försterling, 1989; Hewstone & Jaspars, 1987; Hilton & Slugoski, 1986; Jaspars et al., 1983; Kelley, 1967, 1973; Orvis et al., 1975) and (b) to test whether the causal inference process is unbiased when the relevant set of information is considered. To accomplish our goals, we took the strategy of specifying information over the entire cube. By adopting this strategy, we do not mean to imply that people typically have complete information. It is almost certainly the case that they typically do not. Rather, our position is simply that a test of any model of inference must be considered for the set of events and dimensions to which subjects attend. Unless we know this focal set, it will be impossible to determine whether the inference process per se is inherently biased. Because there is evidence that people use information in nonconfigurational parts of the cube in making causal attributions (Hilton & Slugoski, 1986; Pruitt & Insko, 1980), specifying information over the entire cube is one way to manipulate the relevant information and thereby test our hypotheses.⁴ Neither do we wish to imply that people always attend to the dimensions of persons, stimuli, and occasions. The particular dimensions, as well as the particular number of dimensions, in our materials were chosen to (a) allow a comparison between our model and previous models and (b) provide an explanation for the apparent biases previously reported for materials incorporating these dimensions.

To contrast our model with previous proposals, we constructed sets of problems for which all problems in a set had the same configuration, but every problem in the set had a different pattern of information over the rest of the cube. All previous models (Hilton & Slugoski, 1986; Jaspars et al., 1983; Orvis et al., 1975), with the exception of Försterling (1989), were configuration based. Hilton and Slugoski (1986) made two predictions for the HLH configuration (but not for other configurations) depending on whether the scenarios involve "scripted" or "nonscripted" events. Because all of our problems used nonscripted events, one interpretation of all configuration-based models is that they predict the same attribution(s) for all problems sharing a configuration (i.e., in a set). In contrast, our model predicts different attributions for each of the patterns of information within each set.

An alternative set of predictions for configuration-based models, however, may be derived by applying the principles underlying those models to information over the entire cube, separating those principles from the literal predictions previously made. Under this approach, Jaspars et al.'s (1983) model yields the same predictions as Försterling's (1989), because the deterministic covariation principle underlies both models. For each of our sets of problems sharing a configuration, that principle predicts no possible attribution for all, or most, of the patterns. For one set of problems, Hilton and Slugoski's (1986) abnormal conditions focus (ACF) model makes many of the same predictions as our model. For the other sets, however, it makes the prediction of no possible attribution for most of our problems, because the effect to be explained is normal rather than abnormal. A systematic shift in causal attributions across problems within a set for *all* sets, as predicted by our model, will be inconsistent with the derived predictions of any previous model.

We addressed the issue of biases in two ways. First, we counterbalanced information over the three dimensions of the cube so that any bias toward a dimension or an information variable cannot be attributed to asymmetries in the input. Notice that such an evaluation of bias is independent of any particular covariation-based model of causal attribution. Second, we constructed patterns of information over the nonconfigurational parts of the cube that, according to our model, should (together with configurational information) lead to interaction attributions that have been reported but that are not predicted by previous models. Obtaining such attributions would demonstrate that these apparent deviations can, in fact, be explained by a normative covariational model. To our knowledge, no explanation of such apparent deviations has ever been proposed, despite their striking presence in many studies.

To counterbalance information across the three dimensions of the cube, we rotated each of five complete patterns of information (detached from the labeled axes) from one axis to each of the other two axes in turn, creating a total of 15 problems. We pivoted the rotations at the front top left corner to keep the target event constant. (Equivalently, the pattern of information can remain stationary and the axis labels can be moved.) An example of this rotation is shown in Figure 3 for the pattern of information in which one face of the cube is different from the rest of the cells. In Figure 3a, the face of the cube corresponding to Person 1 differs from the rest of the cube (i.e., the effect occurs in that face of the cube but not in the rest of the cube). By rotating this pattern to each of the other two axes, it is possible to construct two other structurally equivalent patterns of information for which the face of the cube representing Stimulus 1 (see Figure 3b) or Occasion 1 (see Figure 3c) differs from the rest of the cube. Notice that the pattern of information in Figure 3a has the LLH (low consensus, low distinctiveness, high consistency) configuration. Rotating a pattern with that configuration results in patterns with the HHH and HLL configurations (see Figures 3b and 3c, respectively). These are three of the four configurations we used in our materials. We also used the HLH configuration.

Because our problems were constructed by rotating patterns of information as described above, our set of problems is sym-

chotte's (1963) classic experiments, for example, did not involve manipulation, and yet induced the perception of causality (when certain temporal and spatial constraints were satisfied). In our view, manipulation plays an indirect role in two ways: by being a means of establishing psychological temporal priority and by eliminating the possibility of increases in the probability of the effect due to other potential causal factors. Our assumption concerning perceived temporal priority is one that we will neither systematically test nor discuss, although it seems that this assumption, along with others such as conditional independence (Reichenbach, 1956), may play an important role in explaining the realization (by at least some people) that no causal conclusions should be drawn from a correlational study, even when the correlation is strong.

⁴ We leave to subsequent work the task of examining the assumptions and causal attributions made in the face of incomplete information.

(a) Pattern of information indicating a strong main-effect contrast for Person 1 (LLH Configuration, Problem 1).



(b) Pattern of information indicating a strong main-effect contrast for Stimulus 1 (HHH Configuration, Problem 4).



(c) Pattern of information indicating a strong main-effect contrast for Occasion 1 (HLL Configuration, Problem 7).



Figure 3. Three orientations of the pattern of information for Structural Equivalence Set A.

metrical across the three dimensions of persons, stimuli, and occasions. Therefore, each of the three main-effect attributions (to the specific person, to the specific stimulus, and to the specific occasion) should occur equally often. Likewise, each of the two-way interaction attributions should occur equally often. Given the completeness and symmetry of information in our problems, any reliable biases obtained would be attributable to the inference process.

Method

Subjects

The subjects were 91 students at the University of California, Los Angeles (UCLA). Of these, 21 participated in partial fulfillment of requirements for their introductory psychology class. The remaining 70 students participated during small-group sections for their human information processing course (required for the psychology major) as part of a class discussion of reasoning. The course did not cover any topic on inductive reasoning before the students completed the experiment. Because these two groups of subjects did not differ in their responses, all analyses were based on the combined set of data.

Subjects participated in groups of 3 to 20 people. The experiment lasted approximately 20-30 min.

Materials

Structure

Six of our problems incorporated the HLH (high consensus, low distinctiveness, high consistency) configuration. Sets of three problems each had the LLH, HHH, and HLL configurations. As discussed earlier, the various problems sharing a configuration differed in the pattern of information specified for the cells in Regions 4–7. The pattern of information for all 15 problems is given in Table 1. For each region of the cube, a plus indicates that the effect of interest occurs, whereas a minus indicates that the effect does not occur. The cell corresponding to Person 1, Stimulus 1, and Occasion 1 (i.e., Region 0) contains a plus for all problems. Problems that are structurally equivalent, as defined earlier, are labeled with the same letter in the third column of Table 1. The five structurally distinct patterns of information are illustrated in Figures 4 (Sets A, B, and C), 5 (Set D), and 6 (Set E). All problems are presented in Appendix A.

The LLH (low consensus, low distinctiveness, high consistency), HHH, and HLL problems (Nos. 1–9) were constructed by rotating each of three distinct patterns of information to create three sets of three structurally equivalent problems each. The three information patterns are shown in Figure 4, illustrated for the LLH configuration (Problems 1–3). A plus in the figure indicates the presence of the effect, and a minus indicates the absence of the effect. The LLH configurational in-

(a) Pattern of information specified in Problem 1 (LLH; Set A).



(b) Pattern of information specified in Problem 2 (LLH; Set B).



(c) Pattern of information specified in Problem 3 (LLH; Set C).



Figure 4. The patterns of information for Structural Equivalence Sets A, B, and C illustrated using the LLH configuration. (The shaded regions indicate configurational information.)

				Region of the cube							
Problem ^a	Configuration ^b	Structural equivalence set ^c	0	l	2	3	4	5	6	7	
1	LLH	A	+		+	+	-	+	-		
2	LLH	B	+		+	+	+	+	-	+	
3	LLH	C	+		+	+	-	+	+	+	
4	ННН	A	+	+	-	+	+	-	-		
5	ННН	B	+	+		+	+	-	+	+	
6	ННН	C	+	+		+	+	+	-	+	
7	HLL	A	+	+	+	-	-	-	+	-	
8	HLL	B	+	+	+			+	+	+	
9	HLL	C	+	+	+		+	-	+	+	
10 11 12 13 14 15	HLH HLH HLH HLH HLH HLH HLH	D D E E E	+ + + + +	+ + + + +	+ + + + +	+ + + + +	+ - + + +	+ + - + -	- + + - +		

l able l				
Pattern of Information	for Each of	the 15 Expe	rimental I	Problems

Note. + indicates that the effect occurs; - indicates that the effect does not occur.

^a Numbers correspond to problem labels used in the text and the appendixes. ^b The three letters refer, respectively, to high (H) versus low (L) consensus, distinctiveness, and consistency. ^c Problems in the same set specify patterns of information that are structurally equivalent rotations of each other.

formation appears in the shaded regions, and it remains constant across the three problems. The nonconfigurational information in the unshaded regions, however, changes across problems. (Notice that Figure 4a represents the same information as Figure 3a, except that the different occasions are laid out across the page rather than receding in depth, and presence versus absence of the effect is indicated by pluses and minuses rather than by different shading.) For each of the LLH patterns shown in Figure 4, the two rotations illustrated in Figures 3b and 3c created, respectively, corresponding problems with the HHH (Problems 4–6) and HLL (Problems 7–9) configurations.

The HLH problems (Nos. 10–15) were constructed by rotating each of two distinct patterns of information to create two sets of three structurally equivalent problems each. Figure 5 illustrates the three orientations of a pattern in which two faces of the cube differ from the rest of the cells. Figure 6 illustrates those of a pattern in which one face and one beam of the cube differ from the remaining cells. In each figure, the configurational information appears in the shaded regions. Because the HLH configurational information is symmetrical across the three dimensions of the cube (unlike the information specifying the LLH, HHH, and HLL configurations), rotating a pattern containing the HLH configuration in the unshaded regions differs across all six problems in the two figures.

To summarize, the 15 problems were constructed from three orientations of each of five structurally distinct patterns of information over the cube. Because two of the patterns shared the HLH configuration, there were only four configurations represented in the problem set. Each of the 15 problems specified a different set of information. As our problems clearly illustrate, information in the nonconfigurational parts of the cube can assume a vast number of patterns besides those investigated by previous researchers (Hilton & Slugoski, 1986; Pruitt & Insko, 1980).

Predictions

Table 2 presents the attributions predicted by various models of causal inference for our 15 problems. To assess the principles underly-

ing previous configuration-based models—independent of the information to which those models have previously been applied—we derived predictions for those models, whenever we thought possible, by applying their principles to information over the entire cube as specified in our problems. In the table, multiple alternative causes are separated by commas, multiple necessary causes are linked by a multiplication sign, inhibitory causes are denoted by a horizontal bar above the attribution, and weak causes are enclosed in parentheses. We distinguish the ambiguous circumstance attribution (denoted by C) from the unambiguous occasion attribution (denoted by O).

Template model. To apply the model of Orvis et al. (1975) to our complete information problems, it seems that new templates, and perhaps new rules for applying those templates, would be needed, reflecting the fact that information other than the configuration may be relevant for causal attribution. Unfortunately, Orvis et al. did not explicitly state any principles that could be used to derive these new templates. We therefore failed to extend their model, and instead list its predictions for our problems (in the column of Table 2 labeled Template model) based solely on the configuration represented in each problem.

Deterministic covariation principle. Jaspars et al.'s (1983; also see Hewstone & Jaspars, 1987) inductive logic model was designed to apply to configurational information, with the realization that the information is incomplete. However, the deterministic covariation principle underlying their model (identical to Försterling's model, 1989) can be applied to information over the entire cube. This interpretation of the covariation principle states that a factor, or conjunction of factors, is designated a cause if and only if it is both necessary and sufficient for the effect to occur. We applied the deterministic covariation principle to our problems and list its predictions in the corresponding column in Table 2. As an examination of the five structurally distinct patterns illustrated in Figures 4-6 reveals, for only one of the patterns (see Figure 4a, Structural Equivalence Set A) is there a factor that is both necessary and sufficient for the occurrence of the effect. For the remaining four patterns (Structural Equivalence Sets B-E), no factor or conjunction of factors is both necessary and sufficient. The deterministic covariation (a) Pattern of information specified in Problem 10 (HLH; Set D).



(b) Pattern of information specified in Problem 11 (HLH; Set D).



(c) Pattern of information specified in Problem 12 (HLH; Set D).



Figure 5. Three orientations of the pattern of information for Structural Equivalence Set D. (All orientations have the HLH configuration. The shaded regions indicate configurational information.)

principle therefore predicts attributions for only 3 of the 15 problems (Nos. 1, 4, and 7).

We now illustrate the deterministic covariation principle's prediction of no possible causal attributions for two problems from two different structural equivalence sets. Consider a situation in which each of two factors increases the likelihood of the target effect, as in Problem 10 (Figure 5a, Set D). According to our model, Person 1 and Stimulus 1 in that problem are each a cause of the target effect. According to Försterling (1989, p. 621), however, "subjects would attribute an event as being caused by the person if the event was present for the person at all times and at all tasks and when the event was not present for other persons at the same and other tasks at all times." The first condition is met by the pattern of information in Figure 5a but the second is not, because the effect is present for other people (namely, when they encounter Stimulus 1). In other words, because neither Person 1 nor Stimulus 1 is necessary (i.e., the effect can occur in the absence of one or the other of these two factors), the deterministic principle predicts that neither factor should be identified as a cause. Moreover, no other factor or conjunction of factors is both necessary and sufficient, so the deterministic covariation principle predicts that no causal attribution is possible.

Now consider a situation in which our model predicts an interaction contrast, as in Problem 2 (Figure 4b, Set B). For this problem, our model predicts an interaction contrast of Person 1 \times Occasion 1. According to Försterling (1989, p. 621), however, subjects "would attribute an effect to an interaction of two factors when it only occurred when these two causes were present and if it would not occur in the absence of one of the causes." But the conjunction of Person 1 and Occasion 1 (an attribution of something special about this person on this occasion) fails to meet these more stringent criteria demanded by the deterministic covariation principle. It is clearly not the case that the effect *only* occurs when Person 1 and Occasion 1 are present. Because no other factor or conjunction of factors is both necessary and sufficient for the occurrence of the effect, the deterministic covariation principle predicts that no causal attribution is possible.

Abnormal conditions focus (ACF) model. According to Hilton and Slugoski, commonsense reasoning is based on "the counterfactual and contrastive criteria of causal ascription, as unified in the notion of an abnormal condition" (p. 76), rather than on the covariational criterion proposed by others. The counterfactual criterion (consideration of a case in which a possible cause is absent) is used to determine whether a particular factor is necessary for the occurrence of the effect. As Hilton and Slugoski pointed out, the counterfactual criterion by itself cannot account for causal induction, because typically there are many conditions necessary for the occurrence of any given effect. For example (from Hart & Honoré, 1959), the speed of a train, the presence of a faulty rail, and the weight of the cars may all have been necessary for a particular train derailment, but ordinarily only the faulty rail will be indicated as the cause of the derailment (the other necessary factors being relegated to the status of mere conditions). Borrowing from the philosophical literature (Hart & Honoré, 1959; Mackie, 1974), Hilton

(a) Pattern of information specified in Problem 13 (HLH; Set E).



(b) Pattern of information specified in Problem 14 (HLH; Set E).



(c) Pattern of information specified in Problem 15 (HLH; Set E).



Figure 6. Three orientations of the pattern of information for Structural Equivalence Set E. (All orientations have the HLH configuration. The shaded regions indicate configurational information.)

Problem*	Configuration ^b	Template model	Deterministic covariation principle	Abnormal conditions focus model	Probabilistic contrast model
1	LLH	Р	P	Р	Р
2	LLH	P	nc	nc	$P \times O, \overline{O}, (P)$
3	LLH	P	nc	nc	$P \times S, \overline{S}, (P)$
4	ннн	s	s	\$	s
5	ннн	Š	nc	nc	$\mathbf{P} \times \mathbf{S} \stackrel{\mathbf{P}}{\mathbf{P}} (\mathbf{S})$
6	ннн	S	nc	nc	$S \times O, \overline{O}, (S)$
7	HLL	P.S.C	0	0	0
8	HLL	P. S. C	nc	nc	$S \times O, \overline{S}, (O)$
9	HLL	P, S, C	nc	nc	$P \times O, \overline{P}, (O)$
10	нн	PS	nc	PS	PSPXS
10	нин	P S	nc	P.O	POPYO
12	нин	P S	nc	S O	$SO\overline{SXO}$
13	нін	P.S	BC	Р, С	$P, S \times O, (S, O, \overline{P \times S}, \overline{P \times O})$
14	HLH	P. S	nc	ŝ	S. $P \times O$. (P. O. $\overline{P \times S}$, $\overline{S \times O}$)
15	HLH	P, S	nc	õ	$O, P \times S, (P, S, \overline{P \times O}, \overline{S \times O})$

1 4010 2			
Predictions Made by Various	Models for the	Problems Use	d in the Experiment

Note. In columns 3-6, P = person, S = stimulus, O = occasion, C = circumstance, and nc = no cause. Multiple alternative causes are separated by commas, and multiple necessary causes are linked by \times . Inhibitory factors are denoted by a bar above an attribution. Weak causes are enclosed in parentheses. ^a Numbers correspond to problem labels used in the text and the appendixes. ^b The three letters refer, respectively, to high (H) versus low (L) consensus, distinctiveness, and consistency.

and Slugoski (1986) introduced their contrastive criterion of causal selection: "Abnormal conditions come to be dignified as the cause of an event because they are the necessary conditions for the occurrence of a target event that contrast with the conditions obtained in a comparison case where the target event did not occur" (p. 77). In the case of the train derailment, the faulty rail is abnormal, because train rails typically are not faulty. In contrast, the speed of the train, for example, is presumably normal.

Table 2

In applying the notion of abnormality to configurational stimuli, however, Hilton and Slugoski (1986) made a subtle, but important, change in what it means for something to be abnormal. They suggested, for example, that high distinctiveness information (the target person hardly ever exhibits the target behavior in response to other stimuli) serves to "throw the target stimulus into focus as abnormal" (p. 77). Unlike the faulty rail, which was abnormal in the sense of occurring with low probability, the target stimulus is not abnormal in that it per se occurs with lower probability than other conditions (the target person and the target occasion). Rather, it is abnormal in the sense of being associated with an abnormal target effect, which is what occurs with low probability. Nevertheless, this new sense of the contrastive criterion could serve as a model of causal induction if the notion of association with an abnormal effect were clearly defined. Unfortunately, no explicit definition of the concept was provided. From the particular predictions of the ACF model, however, it seems reasonable to us to infer that a factor is associated with an abnormal target effect if that effect occurs often in its presence and rarely in its absence. This interpretation implies that despite Hilton and Slugoski's treatment of their model as an alternative to covariation, it is, in fact, a special case of covariationthe case in which the target effect is abnormal.

Our predictions for the ACF model listed in Table 2 are based on the above interpretation of Hilton and Slugoski's (1986) model. We applied our interpretation of "association with an abnormal effect" to each potential causal factor for each of our problems.⁵ For Problems 1, 4, 7, and 10–15, the effect to be explained is abnormal. For four of these

problems (Nos. 1, 4, 7, and 10), the predictions we derived correspond to those explicitly specified by the ACF model for nonscripted problems with the corresponding configurational information. (Our other problems specified patterns of information for which Hilton and Slugoski did not make any explicit predictions.) Problems 2-3, 5-6, and 8-9 provide an interesting test of the ACF model. For these six problems, the effect to be explained is normal (see Figures 4b and 4c for the two patterns of information for these problems). More formally, if we let the number of values on each of the three dimensions in these problems be n, then the effect occurs in $[n(n-1) + 1]/n^2$ of the cells, whereas it fails to occur in only $(n-1)/n^2$ of the cells (i.e., the effect occurs approximately n times as often as it fails to occur). According to the ACF model, only abnormal factors or factors associated with abnormal effects are dignified as causes. It therefore predicts no possible attribution for these six problems (see Cheng & Novick, 1989b, for a more detailed exposition of this point). Note that although the ACF model makes the same prediction for these problems as does the deterministic covariation principle, it does so for a very different reason.

⁵ Additional assumptions are required for our interpretation of "association with an abnormal effect" to apply sensibly to the conjunction of causal factors, for the following reason. If a causal factor is associated with an abnormal effect, then conjunctions of this factor with some other potential factors are necessarily also associated with the effect. In visual terms, if the effect occurs with higher frequency for a face of the cube (relative to the rest of it), it also occurs with higher frequency relative to the rest of the cube for at least some rows and columns of that face (corresponding to conjunctions of two factors) and for at least some cells of that face (corresponding to conjunction of causal factors, our interpretation of association makes the undesirable prediction that an attribution to a simple factor is always accompanied by attributions to some conjunctions of that factor with other factors.

Probabilistic contrast model. Finally, the predictions of our probabilistic contrast model are listed in the last column of Table 2. Note that in contrast to all previous models, our model predicts a different set of causal attributions for each problem because of the problems' differing patterns of information over the entire cube. Also note that the HLH configuration, previously assumed to lead to person and stimulus attributions or to no attributions at all (e.g., Jaspars et al., 1983; Orvis et al., 1975), can, indeed, lead to any main-effect or interaction attribution, given the appropriate pattern of information over the entire cube.

We derive the predictions of our model in Appendix B for five problems, each representing one of the five distinct patterns of information: LLH Problems 1-3 (see Figures 4a-4c, respectively) and HLH Problems 12 and 15 (see Figures 5c and 6c, respectively). Predictions for the remaining problems can be obtained simply by switching the axis labels. All predictions were derived algebraically, assuming that there are n values on each of the three dimensions.

We illustrate here the application of our model to Problem 12, for which it predicts two facilitatory main-effect causes and an inhibitory interaction cause. As Figure 5c shows, the effect occurs in the left (Stimulus 1) and front (Occasion 1) planes of the cube, but not elsewhere. Clearly, the proportion of cells in which the effect occurs is higher on Occasion 1 (P[occasion-1] = 1) than on other occasions (P[\sim occasion-1 = 1/n. Likewise, it is higher for Stimulus 1 than for other stimuli. However, it is no higher for Person 1 (the top plane) than for other people (the rest of the cube). Our model therefore predicts stimulus and occasion (but not person) as alternative causes. Now consider the Stimulus $1 \times$ Occasion 1 interaction contrast. When Stimulus 1 is present (the left plane), there is no difference in the proportion of cells in which the effect occurs for Occasion 1 (P[stimulus-1, occasion-1] = 1) as opposed to other occasions (**P**[stimulus-1, \sim occasion-1] = 1). In contrast, when Stimulus 1 is absent (all areas of the cube except the left plane), there is a large positive difference in the proportion of cells in which the effect occurs for Occasion 1 ($P[\sim stimulus-1, occasion-1] = 1$) versus other occasions ($P[\sim stimulus-1, \sim occasion-1] = 0$). Therefore, the interaction contrast, being in this case the difference between "no difference" (the contrast involving Occasion 1 and other occasions in the presence of Stimulus 1) and a large "positive difference" (the contrast involving Occasion 1 and other occasions in the absence of Stimulus 1), is a large negative difference. Thus, the interaction is inhibitory. Applying the model similarly to the other interaction contrasts will reveal that both are zero.

Presentation

To ensure that subjects' responses reflected inferences based on the experimental materials rather than on knowledge retrieved from memory, we constructed problems that involved fictional or unfamiliar events. To reduce possible conflict with prior knowledge, subjects were told that all problems concerned events in an imaginary land, where customs and preferences often differ substantially from ours. Because each problem specifies a large amount of novel information, which is not typical in everyday situations, the extra memory load imposed is likely to lower performance. Therefore, our results are likely to underestimate performance in everyday inference.

Subjects were asked to explain what caused a target event. The target event, which was printed in italics, was the first sentence in each problem. The rest of the problem contained two to four sentences, which were numbered and printed in normal font. For example, Problem 6 (HHH configuration) was as follows:

Subjects were told that the information presented after the italicized

event would help them determine its cause. Following the numbered sentences was a list of seven responses. Subjects were told to check "one or more" responses and to check "the minimum number of responses that still give a complete explanation" of the italicized event. The first three responses corresponded to main-effect contrasts of person, stimulus, and occasion, for example, "There is something special about this occasion in general." The next three responses corresponded to interaction contrasts involving two of the three dimensions, for example, "There is something special about the combination of Beth and the morning prayer (only when they are together)." The final response corresponded to an interaction contrast involving all three dimensions, for example, "There is something special about the combination of Beth and the morning prayer and this occasion (only when they are together)." The format for these responses was adapted slightly from that used by Hilton and Slugoski (1986).

To avoid introducing any spurious biases due to variations in wordings across problems, all structurally equivalent problems had similar wordings (see Appendix A). For seven of the problems, the event to be explained was "positive" (e.g., Beth said a particular prayer on this occasion). For the remaining eight problems, the target event was "negative" (e.g., Wendy did not like a particular drink on this occasion). Hilton and Jaspars (1987) explicitly compared positive and negative target events, finding essentially the same attributions in both cases. A different content domain was used for each problem. (To better equate the problems within a configuration, we held the content domain constant across problems sharing a configuration in a subsequent replication, which we report in a later section.)

The 15 problems were distributed among four booklets of four problems each. The first problem in each booklet was one of the three problems for which there is general agreement that the correct attribution is a single main-effect contrast (Problems 1 [LLH], 4 [HHH], and 7 [HLL]). These were expected to be the easiest problems. For each booklet, these three problems appeared approximately equally often as the first page. The remaining three problems for each booklet were fixed. Their order, however, was counterbalanced across subjects using a Latin-square design. These problems were assigned to booklets so as to maximize the diversity of predicted responses within each booklet. The four booklets contained Problems (a) 2, 12, and 15; (b) 3, 6, and 14; (c) 5, 11, and 13; and (d) 8, 9, and 10. The booklets were randomly assigned to subjects.

Results

Scoring the Data

For each problem, subjects could check as many causal attributions as they thought were appropriate. One way to analyze the data would be to compute for each problem the percentage of subjects choosing each attribution. This scheme seems appropriate for addressing the issue of whether subjects have a bias to make certain attributions more than others, regardless of our predictions. We will therefore present our data using this scoring method for the purpose of examining biases.

With the above method, however, the percentage of subjects who chose a predicted attribution includes subjects who also

Beth said the morning prayer on this occasion.

^{1.} In fact, everyone said the morning prayer on this occasion.

^{2.} But nobody said any other prayer on this occasion.

^{3.} On all other occasions, everyone said all the prayers.⁶

⁶ As illustrated in this example, the factors in our materials that yield strong facilitatory contrasts are sufficient to produce the target effect. Such patterns of information were chosen in order to focus on clear-cut contrasts for which the criterion of a noticeable difference is not an issue. Note that except for Problems 1, 4, and 7, these factors were not necessary to produce the effect. Our model requires neither sufficiency nor necessity.

chose unpredicted attributions. A stricter scoring criterion would be to count the number of subjects who chose the predicted attributions and no unpredicted attributions. We will adopt this more conservative scheme to test the differential predictions of our model across problems sharing a configuration. To restrict our analyses to clear-cut contrasts for which the criterion of a noticeable difference is not an issue, we focused on strong contrasts (with values of 1 or approximately 1), which we predict should be included as causes, versus contrasts with value zero, which we predict should be excluded as causes.

For our materials, the interpretation of two types of attributions---those corresponding to weak and inhibitory contrasts--is ambiguous. As can be seen in Appendix B, our problems yield contrasts with absolute values of 1 or approximately 1. 1/n or approximately 1/n, and 0, where n is the assumed number of values on each dimension. It seems debatable whether or not weak contrasts with values of approximately 1/n should be included as causes. The answer depends on the adopted criterion of a noticeable difference and the assumed value of n, both of which are likely to vary from subject to subject. We therefore treat the inclusion of attributions corresponding to such contrasts as optional in our scoring. For a different reason, the interpretation of the selection of inhibitory factors is also ambiguous. Although inhibitory factors (e.g., the presence of a vaccine in one's body) cannot be considered explanations for why an event occurred (e.g., the person contracts the disease), inhibitory as well as facilitatory contrasts are consistent with our response format ("There is something special about . . ."). Therefore, we also treat the inclusion of attributions corresponding to inhibitory contrasts as optional.

In sum, for each problem we considered the following pattern of responses as providing evidence for the probabilistic contrast model: one or more attributions corresponding to strong facilitatory main-effect or interaction contrasts (because each is a cause) and *no* responses corresponding to contrasts of zero. Attributions corresponding to weak or inhibitory causes (as predicted by the respective contrasts) were regarded as optional. Note that subjects who chose *only* inhibitory or weak factors, without at least one strong facilitatory factor, were *not* counted as providing evidence for our model.

Evaluating the Probabilistic Contrast Model

Whereas previous models often predict a single response pattern for all problems sharing a configuration, our probabilistic contrast model predicts a different response pattern for each problem because of the problems' differing patterns of information over the entire cube. In this section, we examine the problems sharing each configuration for evidence of the differential attributions predicted by our model.

The LLH, HHH, and HLL Configurations

As indicated in Table 2, our model predicts a single maineffect contrast for Problems 1 (LLH), 4 (HHH), and 7 (HLL). For each of the remaining two problems for each configuration, our model predicts a different strong facilitatory interaction contrast. It also predicts a strong inhibitory main-effect contrast and a weak facilitatory main-effect contrast.

Table 3

Percentage of Subjects Choosing Various Causal Attributions for the LLH, HHH, and HLL Problems

				Response categories			
Configuration	Problem	n	Р	$P \times O \pm O/P$	$P \times S \pm S/P$	Other	
LLH	1	28	50	11	18	21	
LLH	2	23	9	35	0	57	
	3	25	12	0	80	8	
-				$P \times S +$	S×0+		
			S	P/S	O/S	Other	
ннн	4	30	67	0	10	23	
ННН	5	21	10	76	0	14	
ННН	6	25	0	0	84	16	
-				S×O±	P×O±		
			0	S/O	P/O	Other	
HLL	7	33	45	27	0	27	
HLL	8	22	36	45	0	18	
HLL	9	22	14	0	64	23	

Note. In columns 5-6, / represents and/or, \times links multiple necessary causes, and \pm means with or without; \times has priority over / and \pm , and / has priority over \pm . For example, P \times O \pm O/P means subjects chose the person by occasion interaction attribution; in addition, they may also have chosen the occasion and/or person attributions (but they need not have chosen those attributions). Entries in cells predicted by the probabilistic contrast model are in boldface. The three letters in the first column refer, respectively, to high (H) versus low (L) consensus, distinctiveness, and consistency.

Table 3 lists the percentages of subjects choosing various causal attributions for the sets of problems with the LLH, HHH, and HLL configurations. The number of subjects who received each problem is provided in the column labeled n. The response categories in the table indicate the percentages of subjects who made only attributions predicted for each problem, as defined by the stricter scoring scheme discussed earlier. For each configuration, the four response columns, respectively, list the percentage of subjects who chose only (a) the strong maineffect contrast predicted for the first problem in the set; (b) the strong interaction contrast predicted for the second problem in the set, with or without the appropriate weak and/or inhibitory factors; (c) the strong interaction contrast predicted for the final problem in the set, with or without the appropriate weak and/ or inhibitory factors; and (d) all other response patterns. (The category of other responses includes subjects who chose (a) the predicted strong facilitatory attribution plus one or more attributions corresponding to contrasts of zero or (b) the inhibitory and/or weak facilitatory factors but not any strong facilitatory factors.) Percentages sum to 100 for each row.

The data for each configuration were analyzed separately by comparing the three problems in terms of the frequency of responses falling into each of the three predicted categories.⁷ As

⁷ We excluded "other" responses from all of the statistical analyses testing the ability of our model to discriminate attributions across prob-

is evident from Table 3, responses to problems sharing a configuration varied systematically depending on the pattern of covariation in the entire cube: $\chi^2(4, N = 55) = 47.9, p < .001$, for the LLH problems; $\chi^2(4, N = 62) = 87.8$, p < .001, for the HHH problems; and $\chi^2(4, N = 59) = 47.9, p < .001$, for the HLL problems. This strong support for our model is not due to the inclusion of subjects who chose an inhibitory factor (accompanied by the predicted facilitatory factor), as the pattern of results remains the same if those responses are excluded, with $\chi^2(4, N = 47) = 36.8, p < .001$, for the weakest comparison. For each response category listed in Table 3, the problem for which that response pattern was predicted showed a higher percentage of responses than did the other two problems with the same configuration for which it was not predicted. Each of the nine (three response categories for each of three problem sets) follow-up $\chi^2(1)$ analyses of response patterns (predicted or not) by problems (whether or not the problem was the one for which the pattern was predicted) was significant at least at the .02 level. Clearly, responses to problems with the same configuration varied systematically depending on the pattern of covariation over the entire cube.

Across the three problems with the LLH, HHH, and HLL configurations, an average of 55%, 76%, and 52% of the subjects, respectively, chose only attributions predicted by our model, with an overall average of 61% across configurations. As mentioned earlier, our measure for evaluating our model is quite stringent. An indication of the conservativeness of our measure may be obtained by comparing the more stringent and more lenient scoring criteria for the three problems (1, 4, and 7) for which there is general consensus across models in predicting a single main-effect attribution (these are problems for which a single factor is both necessary and sufficient to produce the effect). For these problems, an average of 54% of the subjects chose the predicted main-effect attribution but not any attributions corresponding to contrasts of zero (see Table 3), whereas an average of 80% of the subjects chose the predicted attributions but also included one or more other responses (see Table 6).

The HLH Configuration

For our analyses, the problems sharing the high consensus, low distinctiveness, high consistency configuration were divided into two subsets: (a) Problems 10–12 and (b) Problems 13–15. The three patterns of information represented by the problems within each subset were structurally equivalent rotations of each other.

The results for Problems 10-12 are presented in Table 4. For these problems, our model predicts two strong facilitatory main-effect contrasts and a strong inhibitory interaction contrast (see Table 2). In consecutive sets of three response columns, the table lists the percentages of subjects who chose *only* (a) one main-effect contrast, (b) two main-effect contrasts, and (c) one or both main-effect contrasts together with the corresponding inhibitory interaction contrast. The final column indicates the percentage of all other responses for each problem (e.g., subjects who chose only an interaction attribution). In accord with our scoring scheme explained above, for each problem lem we considered the following pattern of responses as providing evidence for the probabilistic contrast model: either or both of the strong facilitatory main-effect factors, with or without the inhibitory interaction factor, and no responses corresponding to contrasts of zero. Entries in the predicted cells in Table 4 are in boldface.

As is evident from examining the table, responses for the three problems varied systematically as predicted by our model. To show this statistically, we performed three frequency analyses on the data. Each analysis compared a particular problem to the other two problems combined on the frequency of responses predicted for the particular problem versus the responses not predicted for that problem. For example, one analvsis compared Problem 10 versus the other two problems combined on the frequencies of responses that fit the pattern predicted for Problem 10 (response columns P, S, P & S, and P/S & $P \times S$) versus those that fit the patterns predicted for the other two problems (response columns 3, 5, 6, 8, and 9). The analyses focusing on Problems 10, 11, and 12 were highly significant, yielding, respectively, the following results: $\chi^2(1, N = 50) =$ 18.5, p < .001; $\chi^2(1, N = 50) = 15.2$, p < .001; and $\chi^2(1, N = 50) = 15.2$, p < .001; $\chi^2(1, N = 50) = 15.2$, p < .001; $\chi^2(1, N = 50) = 15.2$, p < .001; $\chi^2(1, N = 50) = 15.2$, p < .001; $\chi^2(1, N = 50) = 15.2$, p < .001; $\chi^2(1, N = 50) = 15.2$, p < .001; $\chi^2(1, N = 50) = 15.2$, p < .001; $\chi^2(1, N = 50) = 15.2$, p < .001; $\chi^2(1, N = 50) = 15.2$, p < .001; $\chi^2(1, N = 50) = 15.2$, p < .001; $\chi^2(1, N = 50) = 15.2$, p < .001; $\chi^2(1, N = 50) = 15.2$, p < .001; $\chi^2(1, N = 50) = 15.2$, p < .001; $\chi^2(1, N = 50) = 15.2$, p < .001; $\chi^2(1, N = 50) = 15.2$, p < .001; $\chi^2(1, N = 50) = 15.2$, p < .001; $\chi^2(1, N = 50) = 15.2$, $\chi^2(1, N = 50) = 1$ 50 = 26.1, p < .001. Again, this strong support for our model is not due to the inclusion of inhibitory factors (when accompanied by predicted facilitatory factors), with $\chi^2(1, N = 41) =$ 12.6, p < .001, for the weakest of the three comparisons if those responses are excluded.

In sum, subjects almost never chose attributions that were consistent with the contrasts predicted for another problem in the set but that were not consistent with the contrasts predicted for the problem under consideration. Across the three problems, an average of 73% of the subjects chose only attributions that were consistent with our probabilistic contrast model.

We now turn to the second set of HLH problems (13–15). For these problems (which were particularly difficult), each factor and combination of two factors covaries, at least weakly, with the effect (see Table 2). We clearly cannot use the scoring scheme we have been using for the rest of our problems, because the predicted response categories for the three problems would be highly overlapping (e.g., a response of P and P \times O could be placed in the predicted category for either Problem 13 or 14). To avoid overlapping response categories, we used the criterion of choosing *only* the strong facilitatory factor(s) predicted for a problem.

For each problem, our model predicts a strong facilitatory main-effect contrast and a strong facilitatory interaction contrast involving the other two dimensions. The percentages of subjects choosing one or both of these attributions (and no other attributions) are presented in Table 5. For example, the column labeled P/S \times O indicates the percentage of subjects for each problem whose attributions matched the predictions for Problem 13: Subjects in this column chose only the person attribution and/or the stimulus by occasion interaction attribution

lems. None of the models predict any changes across problems within a configuration set for any of the responses in this category. Moreover, to the extent that the percentage of unpredicted attributions differs across problems, inclusion of those responses in the analyses would spuriously increase the observed attributional difference among the problems. There was little if any systematicity in the attributions relegated to the "other" column for any of the 15 problems (see Table 6).

<i>joi me i min o</i>		1700										
							Res	oonse ca	tegories			
Configuration	Problem	n	Р	S	0	P&S	P&O	S&O	P/S& P × S	P/O& P×O	S/O& S×O	Other
HLH	10	22	23	14	0	36	0	0	5	0	5	18
HLH	11	21	24	0	14	0	10	0	0	10	5	38
HLH	12	23	0	22	26	0	0	17	0	0	17	17

Percentage of Subjects Choosing Various Causal Attributions
for the First Set of HLH Problems

Note. In columns 7-12 (beginning with the column labeled P&S), & represents and, / represents and/or, and \times links multiple necessary causes; \times has priority over / and &, and / has priority over &. For example, P/S & P \times S means that subjects chose the person (P) attribution and/or the stimulus (S) attribution, and they also chose the Person \times Stimulus interaction attribution. Entries in cells predicted by the probabilistic contrast model are in boldface. O = occasion; HLH = high consensus, low distinctiveness, and high consistency.

(each has a high contrast). The next two columns similarly indicate attributions corresponding to strong facilitatory contrasts for Problems 14 and 15, respectively. Entries in the cells predicted by the probabilistic contrast model are in boldface.

Table /

As predicted, responses to Problems 13–15 varied systematically depending on the pattern of covariation in the entire cube, $\chi^2(4, N = 27) = 21.6, p < .001$. As for the first subset of HLH problems, for each of the sets of attributions that was predicted for a particular problem, the problem for which those attributions were predicted showed a higher percentage of responses than did the two problems (combined) for which they were not predicted: All three $\chi^2(1)$ analyses were significant at least at the .03 level.

Notice that the boldface entries in Table 5 do not reflect the percentage of subjects who were consistent with our model, as defined earlier, because the "other" category includes many who chose attributions corresponding to the predicted strong facilitatory contrast(s), together with attributions corresponding to weak or inhibitory contrasts, which have been regarded as optional. Applying the criterion used in our other problems, the percentages of subjects who chose the predicted strong facilitatory contrast(s), with or without a weak and/or inhibitory contrast, and no other responses were 57%, 76%, and 74%, respectively, for Problems 13, 14, and 15.

Table 5

Percentage of Subjects Choosing Various Causal Attributions for the Second Set of HLH Problems

_				Response c	ategories	
Configuration	Problem	n	$P/S \times O$	S/P×O	$O/P \times S$	Other
HLH	13	21	29	5	14	52
HLH	14	25	0	28	0	72
HLH	15	23	13	4	26	57

Note. In columns 4–6, / represents and/or, and \times links multiple necessary causes; \times has priority over /. For example, P/S \times O means that subjects chose the person (P) attribution and/or the Stimulus (S) \times Occasion (O) interaction attribution (and no other attributions). Entries in cells predicted by the probabilistic contrast model are in boldface. HLH = high consensus, low distinctiveness, and high consistency.

In sum, for every configuration we tested, the various problems sharing a configuration differed significantly in the causal attributions they elicited. Moreoever, the problem for which a particular pattern was predicted by our probabilistic contrast model showed a significantly higher percentage of responses consistent with that pattern than did problems with the same configuration for which that pattern was not predicted. These results provide strong support for the use of covariation as measured by probabilistic contrasts in determining causal inferences.

Comparing the Probabilistic Contrast Model With Competing Models

In this section, we specifically compare our probabilistic contrast model with each of the models whose predictions are listed in Table 2: the template model (Orvis et al., 1975), the deterministic covariation principle that underlies the models of Jaspars et al. (1983; Hewstone & Jaspars, 1987) and Försterling (1989), and the abnormal conditions focus model (Hilton & Slugoski, 1986).

Template model. Because the template model predicts a single response or set of responses across problems sharing a configuration, it clearly fails to account for the systematic shifts in responses for such problems for every configuration reported above.

Deterministic covariation principle. The deterministic covariation principle also fails to account for the shifts in responses across problems sharing a configuration. For the six problems with the HLH configuration, this principle predicts no possible attribution throughout. For two of the three problems with the other three configurations, this principle likewise predicts no possible attribution. Subjects' responses systematically shifted across the above 12 problems, as predicted by our probabilistic contrast model.

Readers may have noticed that we did not include the alternative of "no cause" in our response forms. One might assume that if subjects felt that no causal attribution was possible, they would respond randomly in the absence of this response alternative. Under this assumption, our data clearly disconfirm the deterministic principle. Responses to the problems for which that model predicts no possible attribution were far from random, as already reported. Nevertheless, we would have even stronger evidence against that principle if we could show that subjects failed to choose the "no cause" alternative when it was presented as an option.

We therefore conducted a replication experiment on 72 UCLA summer students, who were paid \$5 for their participation. Half of the subjects received the previous response forms and half received revised forms that added the "no cause" alternative (e.g., "There is nothing special about Wendy, the sunshine prayer, or this occasion, or any combination of the three"). Information patterns for problems 1-12 were used, giving us three problems for each of the four configurations. Each subject received one problem from each configuration, with problem order counterbalanced across subjects. To better equate the problems within a configuration, we held the content domain constant across problems sharing a configuration. Also, all problems involved positive events.

Across the 12 problems, only 7.6% of the subjects who received the response form with the "nothing special" alternative chose that response (and no other responses). The deterministic covariation principle underlying the models of Försterling (1989) and Jaspars (Hewstone & Jaspars, 1987; Jaspars et al., 1983) predicts that no attribution is possible for 9 of the 12 problems (all problems except 1, 4, and 7). The low overall percentage of such responses does not support the principle. A more critical question, though, is whether subjects were more likely to give this response for the nine problems for which the deterministic covariation principle predicts no possible attribution than for the three problems for which the principle predicts a causal attribution. The answer is a resounding "no": 7.4% of the subjects chose the "nothing special" response (and no other responses) for the problems for which it was predicted, and a similar 8.3% of the subjects chose (only) that response for the problems for which it was not predicted. These data thus replicate those of the original experiment in disconfirming the deterministic covariation principle.

The inclusion of the "no cause" alternative had no effect on the relative frequencies of other responses for any of the problems. All configurations were analyzed as described above for the original experiment, and the results replicated those reported earlier.

Abnormal conditions focus model. For all of the problems for which the ACF model and our model make no overlapping attributional predictions, the ACF model predicts no possible attribution. This is the same prediction as that made by the deterministic covariation principle for these problems. Our refutation of that model, just described, applies identically to the ACF model.

The Issue of Bias

We addressed the issue of bias to some extent in the section on evaluating the probabilistic contrast model. By obtaining attributions that were not predicted by other models but that have been found in previous research, we demonstrated that these apparent deviations from the predictions of previous models can, in fact, be explained by a normative covariational model. These deviations include interaction attributions, which have not previously been explained.

We now consider the person bias. For the three problems for which a single factor is both necessary and sufficient to produce the effect (Problems 1, 4, and 7), there is no evidence for a bias toward making a person attribution as opposed to a stimulus or occasion attribution: 75%, 80%, and 85% of the subjects chose these respective attributions for these problems (see Table 6). More generally, Table 6 shows, for each problem, the percentage of subjects choosing each of the three main-effect attributions and each of the three interaction attributions, regardless of what other attributions, if any, they chose. The final line reports the average percentage of subjects making each attribution across the 15 problems. Person, stimulus, and occasion attributions were made by 31%, 27%, and 38% of the subjects, respectively. Consistent with our hypothesis of no biases in the inference process, a repeated-measures ANOVA with problems as the replications variable and the three attributions as a withinproblem variable indicated no difference in the prevalence of the three attributions, F(2, 28) < 1. We did not conduct the comparable statistical analysis for the interaction attributions because of the bimodal distribution of percentages for the P \times S and $P \times O$ attributions. However, the results show no obvious biases (see Table 6).

On the basis of the smaller percentage of variance in her data accounted for by consensus information compared with distinctiveness and consistency information, McArthur (1972) concluded that there is a bias against using consensus information. To compare our results with hers, we similarly computed the percentages of variance accounted for by consensus, distinctiveness, and consistency in our data, with the important difference that we redefined these variables in terms of contrasts to capture response variation over the entire cube rather than over only the configuration. For example, consensus was redefined in terms of whether or not there were differences in the target behavior between the person in question and other people on most stimuli across most occasions (rather than in terms of whether or not there were differences across people about a particular stimulus on a particular occasion). Distinctiveness and consistency were redefined similarly. In other words, each of these variables was defined with respect to the same set of events, events in the entire cube. They differed only in the way the cube was sliced. For each (redefined) information variable, we coded three levels of covariation: no covariation, weak facilitatory covariation, and strong facilitatory covariation.

With our redefinitions of consensus, distinctiveness, and consistency, these information variables should be able to predict main-effect attributions. We thus analyzed each of the three dependent variables of person, stimulus, and occasion attributions (regardless of what other attributions, if any, were chosen) in a separate stepwise regression using the redefined information variables as independent variables. (Three equally spaced values were used to code no, weak, and strong covariation.) We excluded from these analyses the six problems for which our model predicts strong inhibitory main-effect contrasts (Problems 2, 3, 5, 6, 8, and 9). Although the patterns of information for these problems indicate high covariation (i.e., low consensus, high distinctiveness, or low consistency for the redefined variables) along a dimension (persons, stimuli, or occasions),

			Attribution						
Problem	Configuration	n	P	S	0	P×S	P×O	\$×0	
1	LLH	28	75	0	14	18	11	7	
2	LLH	23	43	0	48	4	39	22	
3	LLH	25	56	20	0	84	0	0	
4	ННН	30	3	80	10	0	7	13	
5	ННН	21	19	43	0	86	0	5	
6	ннн	25	0	36	52	4	8	88	
7	HLL	33	3	0	85	0	3	33	
8	HLL	22	5	36	77	9	0	45	
9	HLL	22	9	0	50	0	77	0	
10	HLH	22	64	55	0	18	0	9	
11	HLH	21	57	10	38	5	29	29	
12	HLH	23	0	57	61	0	0	35	
13	HLH	21	48	14	48	10	0	24	
14	HLH	25	36	44	32	4	60	16	
15	HLH	23	48	4	52	61	4	35	
М			31	27	38	20	16	24	

Percentage of Subjects Choosing Each of the Three Single-Factor Attributions and Each of the Three Conjunctive-Factor Attributions for Each Problem

Note. Numbers in boldface indicate the percentages of subjects who chose the strong facilitatory attributions predicted by the probabilistic contrast model. Numbers in italics indicate the percentages of subjects who chose the weak facilitatory or (strong or weak) inhibitory attributions predicted by the probabilistic contrast model. The three letters refer, respectively, to high (H) versus low (L) consensus, distinctiveness, and consistency.

the covariation is negative (indicating inhibitory causes). Subjects would therefore be quite justified in not checking the corresponding dependent variable as a response. Thus, each regression considered responses to nine problems (1, 4, 7, and 10-15), with each subject's response to a particular problem as a separate data point in the analysis.

Table 6

These analyses showed that consensus accounted for 31% of the variance in person attributions, distinctiveness accounted for 33% of the variance in stimulus attributions, and consistency accounted for 24% of the variance in occasion attributions. The only independent variable that was significant in each analysis was the one predicted by covariation: F(1, 224) =98.9, p < .001, for consensus predicting person attributions; F(1, 224) = 110, p < .001, for distinctiveness predicting stimulus attributions; and F(1, 224) = 68.7, p < .001, for consistency predicting occasion attributions.

Because the decision of whether to include weak facilitatory contrasts in one's response depends on where one sets the criterion for a noticeable difference, we conducted a separate set of regression analyses on the six problems (1, 4, 7, and 10–12) for which our model predicts strong contrasts only. For this set of problems, consensus accounted for 47% of the variance in person attributions, distinctiveness accounted for 45% of the variance in stimulus attributions, and consistency accounted for 34% of the variance in occasion attributions. Only the predicted independent variable was significant in the analyses of person and stimulus attributions; F(1, 155) = 138, p < .001, for consensus predicting person attributions; and F(1, 155) = 127, p < .001, for distinctiveness predicting stimulus attributions. Consistency accounted for the largest percentage of the variance in occasion attributions.

to enter the stepwise regression, F(1, 155) = 80.3, p < .001. Although the other two independent variables were also significant (F[2, 154] = 7.16, p < .005, for consensus, and F[3, 153] = 8.03, p < .001, for distinctiveness), the additional percentages of variance accounted for were very small (3% for each variable). Thus, in neither subset of problems was there any evidence for a bias against using consensus information.

Discussion

The Probabilistic Contrast Model

We have proposed a model of causal induction that is a probabilistic analogue of statistical contrasts. Rather than requiring the complex quantitative computations underlying statistical contrasts, our model only requires the comparison of proportions. According to our model, causal inference is based on the computation of contrasts between the proportion of times the effect occurs for a particular value on a dimension versus other values on that dimension. In particular, a factor will be designated a cause if the proportion of times the effect occurs when that factor is present is greater (by some criterion) than the proportion of times the effect occurs when the factor is absent. A conjunction of factors (e.g., a particular person in combination with a particular stimulus) will be designated a cause if there is a noticeable difference between (a) the contrast for one of the factors (e.g., the particular person) when the other factor (e.g., the particular stimulus) is present and (b) the contrast for that same factor (the person) when the other factor (the stimulus) is absent.

Like Kelley's (1967, 1973) ANOVA model, the predictions of

our model are based on covariation. Our model, however, is more general than Kelley's. First, covariation is specified by probabilistic contrasts rather than by the information variables of consensus, distinctiveness, and consistency. Probabilistic contrasts can be computed for any dimension. Thus, although we tested the predictions of our model by specifying information over the entire Person \times Stimulus \times Occasion cube described by Kelley, our model is not committed to those particular dimensions. For example, people may attribute effects to countries, as when the presence of giant pandas in China alone is attributed to something special about China, perhaps its climate and vegetation. Neither is our model committed to a cube. The number of relevant dimensions may differ across problem contexts as well as across people. For example, some people may consider the presence or absence of other animals in China and other countries (the stimulus dimension in Kelley's cube) to be irrelevant to the determination of the cause of the exclusive presence of giant pandas in China. Our main thesis is simply that causal inferences are based on the computation of probabilistic contrasts for a focal set-the set of events considered relevant by the attributor.

Demonstrating the generality of our probabilistic contrast model, we (Cheng & Novick, 1989b) have shown that this model, when applied to the set of events and dimensions to which attention is focused, as determined by the pragmatic context (also see Cheng & Holyoak, 1985; Cheng, Holyoak, Nisbett, & Oliver, 1986), can account for the intuitive distinction people draw between causes and conditions that merely enable a cause to realize its effect but are not themselves causes. For example, people are unlikely to say that the presence of oxygen was the cause of a forest fire. Rather, they are likely to reserve the title "cause" for factors such as lightning, a dropped cigarette, or the unusual dryness of the climate. This distinction between causes (e.g., lightning) and conditions (e.g., oxygen) is made despite the knowledge that oxygen and lightning, for example, are individually necessary for a particular fire to occur, and are jointly sufficient (along with other conditions such as the combustibility of wood) to produce the fire. In a series of experiments, we (Cheng & Novick, 1989a, 1989b) (a) ruled out alternative hypotheses, including the abnormality of the factors, their observability, and assumptions about the state of knowledge of the inquirer; and (b) showed that previously proposed models (Jaspars et al., 1983; Kelley, 1967, 1973; Schustack & Sternberg, 1981; Shaklee & Tucker, 1980; Suppes, 1970, 1984) cannot account for the distinction.

Potential Criticisms of Our Methodology

The results of this experiment clearly support our model against previous ones. Several criticisms, however, might be raised against the validity of our methodology. We consider some of them below.

Because our materials specified strong facilitatory contrasts that were, in fact, sufficient, a question may be raised as to whether sufficiency is a criterion that can indeed account for our results. It clearly cannot. Consider, for example, the Person \times Stimulus contrast. The combination of Person 1 and Stimulus 1 is sufficient to produce the effect for Problems 1– 6 and 10–15. Yet, whereas an average of 77% of our subjects attributed the effect to that combination for Problems 3, 5, and 15, for which a strong facilitatory Person \times Stimulus contrast is predicted by our model, an average of a mere 7% made the same attribution for the rest of these problems, for which that contrast is not predicted. Moreover, much previous evidence demonstrates that people do infer insufficient factors to be causes (e.g., Jaspars, 1983; McArthur, 1972).

Despite our intent to create arbitrary materials, it is possible that they may, in fact, have been nonarbitrary. If so, then two alternative objections might be raised against our results. First, if subjects' real-life knowledge about a content domain agreed with our specifications, their responses might have been based on knowledge retrieval rather than on inference. Second, if their knowledge conflicted with our specifications, it might be argued that our methodology is rendered invalid.

We do not think the retrieval of prior knowledge can account for our results. First, it is not at all obvious how our predictions follow from real-life knowledge. For example, given one's reallife knowledge relevant to the statement "Dave would not eat rabbit meat on this occasion," it is not clear that the conclusion "There is something special about the occasion" would follow. Yet 85% of our subjects who received this problem chose the occasion attribution (compared to 0% for the stimulus attribution and 3% for the person attribution), as predicted by the application of our model to the pattern of covariation specified in the problem. Moreover, in our replication, content domain was kept constant across problems sharing a configuration. Yet systematic shifts in attributions predicted by probabilistic contrasts nonetheless occurred. Such shifts clearly cannot be explained by subjects' prior knowledge about a particular domain.

Neither do we believe that possible conflicts between subjects' real-world knowledge and the information specified in our problems render our methodology invalid. Almost certainly, we were not completely successful in manipulating the information used by the subject. Information overload, as we mentioned earlier, and conflict with prior knowledge are both likely to have interfered with our intent. These may, in fact, be important explanations for the residual noise in our data. If we had been more successful in manipulating subjects' assumptions, their attributions might have been even more consistent with the predictions of our model. The partial ineffectiveness of our methodology, however, does not invalidate it. The systematic variation across problems that we found vindicates the effectiveness of our manipulation. Failure to manipulate subjects' assumptions could only explain the potential failure of experiments such as ours-it cannot explain their actual success.

A further possible objection might be based on the possibility that our materials were rather artificial. The artificiality of our materials might have led subjects to devise artificial methods of inference that they do not use when confronted with real-life problems. It seems highly implausible to us, however, that artificial methods spontaneously devised by subjects should systematically support our model. One would expect such methods to be varied and therefore to lead to confusing data. It is particularly unlikely that most subjects would happen to devise a rule that *coincides* with our definition of interaction contrasts.

More important, as mentioned earlier, we (Cheng & Novick, 1989a, 1989b) have shown that the probabilistic contrast

model, but not other models of causal induction, can account for the intuitive distinction people draw between causes and enabling conditions. The materials used in those experiments involved common physical and biological events such as fires and plant growth. Thus, there is evidence that our model explains an intuitive distinction based on nonartificial materials.

Could our subjects have successfully applied a set of trained rules that is equivalent or similar to our model? According to a questionnaire administered at the end of the experiment, 81% of our subjects answered "no" when asked "Have you ever taken any courses in high school or college that included a discussion of causation (i.e., how to determine what are possible causes of events)?" and asked to list the relevant courses. Only 1% listed statistics. (Eight percent listed a psychology research methods laboratory course, 8% listed a philosophy or logic course, and 2% listed other courses.) Even those who have taken a course that included a discussion of causation are unlikely to have been trained on conjunctive causal factors. It is therefore unlikely that our subjects were applying trained rules.

Reassessment of Biases

By specifying for the subject types of information that typically have been ignored in previous experiments on causal attribution and by redefining the information variables of consensus, distinctiveness, and consistency over the entire cube, we found neither a bias for attributing an effect to a person nor a bias against using consensus information. Moreover, by manipulating information in the nonconfigurational part of the cube, while keeping configurational information constant, we were able to elicit various causal attributions, including ones that have previously been regarded as deviations from normative covariation. Our results therefore show that such apparent deviations may be explained by the pattern of the information over which covariation is computed. Just as we were able to vary the information given for the nonconfigurational part of the cube, subjects in previous experiments (e.g., Försterling, 1989; Hilton & Slugoski, 1986; Jaspars et al., 1983; McArthur, 1972; Orvis et al., 1975) might also have varied their assumptions regarding the unspecified parts of the cube from situation to situation and problem to problem, thus producing what appeared to be capricious biases.

Our results also show that the deterministic covariation principle (Försterling, 1989; Jaspars et al., 1983), even when applied to information over the entire cube, does not describe the process of causal induction. Contrary to this principle but in accord with our probabilistic contrast model, our subjects consistently attributed effects to factors that were not necessary to produce the effect but that corresponded to strong contrasts. The large deviations from the deterministic covariation principle do not appear to represent irrational biases, because they systematically follow from a less stringent but nonetheless unbiased probabilistic model.

Because our experiment did not present subjects with the opportunity to spontaneously select dimensions out of a larger set, our results tell us nothing about (a) whether the dimensions of persons, stimuli, and occasions are the ones people spontaneously seek out, and, hence, (b) how prior knowledge affects the selection. What our results do indicate, however, is that from the set of selected dimensions, people compute covariation without bias, and the computed pattern of covariation then determines causal inferences.

References

- Alloy, L. B., & Abramson, L. Y. (1979). Judgment of contingency in depressed and nondepressed students: Sadder but wiser? *Journal of Experimental Psychology: General*, 108, 441-485.
- Alloy, L. B., & Tabachnik, N. (1984). Assessment of covariation by humans and animals: The joint influence of prior expectations and current situational information. *Psychological Review*, 91, 112–149.
- Borgida, E. (1978). Scientific deduction—evidence is not necessarily informative. A reply to Wells and Harvey. *Journal of Personality and Social Psychology*, 36, 477–482.
- Borgida, E., & Brekke, N. (1981). The base rate fallacy in attribution and prediction. In J. H. Harvey, W. Ickes, & R. F. Kidd (Eds.), New directions in attribution research (Vol. 3, pp. 63-95). Hillsdale, NJ: Erlbaum.
- Bruner, J. S. (1957). Going beyond the information given. In H. E. Gruber, K. R. Hammond, & R. Jessor (Eds.), *Contemporary approaches to cognition* (pp. 41–69). Cambridge, MA: Harvard University Press.
- Bullock, M., Gelman, R., & Baillargeon, R. (1982). The development of causal reasoning. In W. J. Friedman (Ed.), *The developmental psy*chology of time (pp. 209-254). Orlando, FL: Academic Press.
- Cheng, P. W., & Holyoak, K. J. (1985). Pragmatic reasoning schemas. Cognitive Psychology, 17, 391–416.
- Cheng, P. W., Holyoak, K. J., Nisbett, R. E., & Oliver, L. M. (1986). Pragmatic versus syntactic approaches to training deductive reasoning. *Cognitive Psychology*, 18, 293–328.
- Cheng, P. W., & Novick, L. R. (1989a). Natural causal induction. Manuscript in preparation, Department of Psychology, University of California, Los Angeles.
- Cheng, P. W., & Novick, L. R. (1989b). Causes versus enabling conditions. Unpublished manuscript, Department of Psychology, University of California, Los Angeles.
- Cheng, P. W., & Novick, L. R. (1990). Where is the bias in causal attribution? In K. Gilhooly, M. Keane, R. Logie, & G. Erdos (Eds.), *Lines* of thinking: Reflections on the psychology of thought (Vol. 1, pp. 181– 197). New York: Wiley.
- Försterling, F. (1989). Models of covariation and attribution: How do they relate to the analogy of analysis of variance? *Journal of Personality and Social Psychology*, 57, 615–625.
- Hansen, R. D., & Donoghue, J. M. (1977). The power of consensus: Information derived from one's own and others' behavior. *Journal of Personality and Social Psychology*, 35, 294–302.
- Harris, R. J. (1977). Comprehension of pragmatic implications in advertising. Journal of Applied Psychology, 62, 603–608.
- Hart, H. L., & Honoré, A. M. (1959). Causation in the law. Oxford, England: Clarendon Press.
- Heider, F. (1958). *The psychology of interpersonal relations*. New York: Wiley.
- Henle, M. (1962). On the relationship between logic and thinking. Psychological Review, 69, 366-378.
- Hewstone, M. R. C., & Jaspars, J. M. F. (1983). A re-examination of the roles of consensus, consistency and distinctiveness—Kelley cube revisited. *British Journal of Social Psychology*, 22, 41–50.
- Hewstone, M. R. C., & Jaspars, J. M. F. (1987). Covariation and causal attribution: A logical model of the intuitive analysis of variance. *Jour*nal of Personality and Social Psychology, 53, 663–672.
- Hilton, D. J. (1988). Logic and causal attribution. In D. J. Hilton (Ed.), Contemporary science and natural explanation: Commonsense conceptions of causality (pp. 33–65). Brighton: Harvester Press.

Hilton, D. J. (1990). Conversational processes and causal explanation. *Psychological Bulletin*, 107, 65–81.

- Hilton, D. J., & Jaspars, J. M. F. (1987). The explanation of occurrences and non-occurrences: A test of the inductive logic model of causal attribution. *British Journal of Social Psychology*, 26, 189-201.
- Hilton, D. J., & Slugoski, B. R. (1986). Knowledge-based causal attribution: The abnormal conditions focus model. *Psychological Review*, 93, 75–88.
- Hilton, D. J., Smith, R. H., & Alicke, M. D. (1988). Knowledge-based information acquisition: Norms and the functions of consensus information. *Journal of Personality and Social Psychology*, 55, 530–540.
- Jaspars, J. M. F. (1983). The process of causal attribution in common sense. In M. R. C. Hewstone (Ed.), Attribution theory: Social and functional extensions (pp. 28-44). Oxford, England: Basil Blackwell.
- Jaspars, J. M. F., Hewstone, M. R. C., & Fincham, F. D. (1983). Attribution theory and research: The state of the art. In J. M. F. Jaspars, F. D. Fincham, & M. R. C. Hewstone (Eds.), *Attribution theory: Essays* and experiments (pp. 3-36). Orlando, FL: Academic Press.
- Jenkins, H., & Ward, W. (1965). Judgment of contingency between responses and outcomes. *Psychological Monographs*, 7, 1–17.
- Johnson, M. K., Bransford, J. D., & Solomon, S. K. (1973). Memory for tacit implications of sentences. *Journal of Experimental Psychol*ogy, 98, 203-205.
- Kassin, S. (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: University of Nebraska Press.
- Kelley, H. H. (1973). The processes of causal attribution. American Psychologist, 28, 107–128.
- Kelley, H. H., & Michela, J. L. (1980). Attribution theory and research. Annual Review of Psychology, 31, 457–501.
- Kulik, J. A., & Taylor, S. E. (1980). Premature consensus on consensus? Effects of sample-based versus self-based consensus information. *Journal of Personality and Social Psychology*, 38, 871–878.
- Mackie, J. L. (1974). *The cement of the universe: A study of causation*. Oxford, England: Clarendon Press.

Marr, D. (1982). Vision. San Francisco: W. H. Freeman and Company.

- McArthur, L. Z. (1972). The how and what of why: Some determinants and consequences of causal attribution. *Journal of Personality and Social Psychology*, 22, 171–193.
- McArthur, L. Z. (1976). The lesser influence of consensus than distinctiveness information on causal attributions: A test of the personthing hypothesis. *Journal of Personality and Social Psychology*, 33, 733-742.

Michotte, A. E. (1963). The perception of causality. New York: Basic.

- Mill, J. S. (1973). A system of logic ratiocinative and inductive (8th ed.).
 In J. M. Robson (Ed.), *Collected works of John Stuart Mill* (Vols. 7 & 8). Toronto, Canada: University of Toronto Press. (Original work published 1843)
- Nisbett, R. E., & Borgida, E. (1975). Attribution and the psychology of

prediction. Journal of Personality and Social Psychology, 32, 932-943.

- Nisbett, R. E., Krantz, D. H., Jepson, C., & Kunda, Z. (1983). The use of statistical heuristics in everyday inductive reasoning. *Psychological Review*, 90, 339-363.
- Nisbett, R. E., & Ross, L. (1980). Human inference: Strategies and shortcomings of social judgment. Englewood Cliffs, NJ: Prentice-Hall.
- Orvis, B. R., Cunningham, J. D., & Kelley, H. H. (1975). A closer examination of causal inference: The roles of consensus, distinctiveness, and consistency information. *Journal of Personality and Social Psychology*, 32, 605-616.
- Pruitt, D. J., & Insko, C. A. (1980). Extension of the Kelley attribution model: The role of comparison-object consensus, target-object consensus, distinctiveness, and consistency. *Journal of Personality and Social Psychology*, 39, 39–58.
- Reichenbach, H. (1956). *The direction of time*. Berkeley: University of California Press.
- Robinson, G. H. (1964). Continuous estimation of a time-varying probability. *Ergonomics*, 7, 7–21.
- Ross, L. (1977). The intuitive psychologist and his shortcomings: Distortions in the attribution process. Advances in Experimental Social Psychology, 10, 174–220.
- Ruble, D. N., & Feldman, N. S. (1976). Order of consensus, distinctiveness, and consistency information and causal attribution. *Journal of Personality and Social Psychology*, 34, 930–937.
- Schustack, M. W., & Sternberg, R. J. (1981). Evaluation of evidence in causal inference. Journal of Experimental Psychology: General, 110, 101-120.
- Shaklee, H., & Tucker, D. (1980). A rule analysis of judgments of covariation between events. *Memory & Cognition*, 8, 459-467.
- Shuford, E. H. (1961). Percentage estimation of proportion as a function of element type, exposure time, and task. *Journal of Experimen*tal Psychology, 61, 430-436.
- Suppes, P. (1970). A probabilistic theory of causality. Amsterdam: North-Holland.
- Suppes, P. (1984). Probabilistic metaphysics. Oxford, England: Basil Blackwell.
- Tversky, A., & Kahneman, D. (1982). Evidential impact of base rates. In D. Kahneman, P. Slovic, & A. Tversky (Eds.) Judgment under uncertainty: Heuristics and biases (pp. 153–160). New York: Cambridge University Press.
- Weiner, B., Frieze, I., Kukla, A., Reed, L., Rest, S., & Rosenbaum, R. M. (1971). Perceiving the causes of success and failure. In E. E. Jones, D. E. Kanouse, H. H. Kelley, R. E. Nisbett, S. Valins, & B. Weiner (Eds.), *Attribution: Perceiving the causes of behavior* (pp. 63– 77). Morristown, NJ: General Learning Press.
- Wells, G. L., & Harvey, J. H. (1977). Do people use consensus information in making causal attributions? *Journal of Personality and Social Psychology*, 35, 279–293.
- Zuckerman, M. (1978). Actions and occurrences in Kelley's cube. Journal of Personality and Social Psychology, 36, 647–656.

Appendix A

Causal Attribution Problems

Problem 1 (LLH):

- Jane had fun washing dishes on this occasion.
- 1. In fact, Jane always has fun doing all chores.
- 2. But nobody else ever has fun doing any chores.

Problem 2 (LLH):

- Cathy is not allergic to olive tree pollen on this occasion.
- 1. In fact, Cathy is not allergic to any other kind of tree pollen on this occasion either.
- 2. But everyone else is allergic to all kinds of tree pollen on this occasion.
- 3. On all other occasions, nobody has been allergic to any kind of tree pollen.

Problem 3 (LLH):

Adam thinks that narcissus (a flower) smells nice on this occasion.

- 1. In fact, Adam has always thought that narcissus smells nice.
- 2. But nobody else has ever thought that narcissus smells nice.
- 3. Everyone has always thought that all other flowers smell nice.

Problem 4 (HHH):

Wendy does not like to drink Campari (a kind of liquor) on this occasion.

- 1. In fact, nobody has ever liked to drink Campari.
- 2. But everyone has always liked to drink all other kinds of alcoholic drinks.

Problem 5 (HHH):

- Sam is not afraid of John's raccoon on this occasion.
- 1. In fact, Sam has never been afraid of John's raccoon.
- 2. But Sam has always been afraid of all other raccoons.
- 3. Nobody else has ever been afraid of any raccoon.

Problem 6 (HHH):

Beth said the morning prayer on this occasion.

- 1. In fact, everyone said the morning prayer on this occasion.
- 2. But nobody said any other prayer on this occasion.
- 3. On all other occasions, everyone said all the prayers.

Problem 7 (HLL):

- Dave would not eat rabbit meat on this occasion.
- 1. In fact, nobody would eat any kind of meat on this occasion.
- 2. But on all other occasions, everyone has eaten all kinds of meats.

Problem 8 (HLL):

Vicky is not wearing a gardenia (a flower) on her collar on this occasion.

- 1. In fact, nobody is wearing a gardenia on his or her collar on this occasion.
- 2. But on all other occasions, everyone has worn a gardenia on his or her collar.
- 3. Nobody has ever worn any other kind of flower on his or her collar.

Problem 9 (HLL):

Patty enjoyed playing majong (a gambling game) on this occasion.

1. In fact, Patty enjoyed playing all kinds of gambling games on this occasion.

- 2. But on all other occasions she did not enjoy playing any gambling games.
- 3. Everyone else has always enjoyed playing all kinds of gambling games.

Problem 10 (HLH):

Kim does not enjoy listening to the zither (a musical instrument) on this occasion.

- 1. In fact, Kim has never enjoyed listening to the zither.
- 2. Nobody else has ever enjoyed listening to the zither either.
- But everyone else has always enjoyed listening to all other instruments.
- 4. However, Kim has never enjoyed listening to any instrument.

Problem 11 (HLH):

Alice displayed sculptures made from clay in her home on this occasion.

- 1. In fact, Alice displayed sculptures made from all kinds of materials on this occasion.
- 2. Alice has displayed sculptures made from all kinds of materials on all other occasions too.
- 3. But on all other occasions nobody else has displayed sculptures made from any kind of material.
- 4. However, on this occasion everyone displayed sculptures made from all kinds of materials.

Problem 12 (HLH):

Fred sang "Golden Slumbers" (a children's song) on this occasion.

- 1. In fact, everyone sang "Golden Slumbers" on this occasion.
- 2. Everyone also sang all other children's songs on this occasion.
- 3. But nobody sang any other children's songs on any other occasion.
- 4. However, everyone has always sung "Golden Slumbers."

Problem 13 (HLH):

Eric did not like to dance the tango on this occasion.

- 1. In fact, Eric has never liked to dance any kind of dance.
- But except for this occasion, everyone else has always liked to dance all kinds of dances (including the tango).
- 3. On this occasion, nobody else liked to dance the tango (but like all other occasions, they did like to dance all other dances).

Problem 14 (HLH):

- Susan did not bow to the statue of fire on this occasion.
- 1. In fact, nobody has ever bowed to the statue of fire.
- 2. But except for this occasion, everyone (including Susan) has always bowed to all other statues.
- 3. On this occasion, Susan did not bow to any other statue (but like all other occasions, everyone else bowed to all other statues).

Problem 15 (HLH):

George likes collard greens (a leafy vegetable) on this occasion.

- 1. In fact, everyone likes all kinds of leafy vegetables on this occasion.
- But except for George, on all other occasions nobody has liked any kind of leafy vegetable (including collard greens).
- 3. On all other occasions, George has always liked collard greens (but like everyone else he hasn't liked any other kind of leafy vegetable).

(Appendixes continue on next page)

Appendix B

Predictions of the Probabilistic Contrast Model for Five Problems

We derive below the predictions of the probabilistic contrast model for five problems, each representing one of the five structurally distinct patterns. For each of the three main effects and each of the three twoway interactions, contrasts are derived for each problem. Because our problems do not specify the number of people, stimuli, and occasions, and because our information patterns are symmetric across the three dimensions, we make the simplifying assumption that there are an equal number of values on every dimension, and let the variable *n* represent this number. Spatial labels for some of the columns refer to the Person × Stimulus × Occasion cube as oriented in Figure 2. In the Predicted attribution column, person, stimulus, and occasion attributions are denoted respectively by P, S, and O. Inhibitory contrasts are denoted by a bar above the letter (e.g., \overline{S}). No attribution corresponding to a contrast is denoted by a dash (—).

According to the probabilistic contrast model, positive contrasts predict facilitatory attributions and negative contrasts predict inhibitory attributions. Contrasts of zero predict no attribution. We make no assumption regarding whether contrasts with an absolute value of 1/n or approximately 1/n are above the criterion for a noticeable difference. Attributions corresponding to contrasts of this size are enclosed in parentheses. (These attributions are treated as optional in our scoring.) Contrasts much greater than 1/n are assumed to be above the threshold, and these predict corresponding attributions.

Person Contrast $(P[i] - P[\sim i])$:

Problem	Relative frequence	cy of target event for		Predicted attribution
	Target person (i) (top plane)	Other people (~i) (other planes)	Contrast	
1	1	0	1	Р
2	1	(n-1)/n	1/n	(P)
3	1	(n-1)/n	1/n	(P)
12	$(2n-1)/n^2$	$(2n-1)/n^2$	0	-
15	$(2n-1)/n^2$	1/n	$(n-1)/n^2$	(P)

Stimulus Contrast (P[j] - P[~j]):

Problem	Relative frequence	y of target event for		
	Target stimulus (j) (left plane)	Other stimuli (~j) (other planes)	Contrast	Predicted attribution
1	1/n	1/n	0	
2	$[(n-1)^2 + n]/n^2$	$[(n-1)^2 + n]/n^2$	0	
3	1/n	1	-(n-1)/n	Ī
12	1	1/ <i>n</i>	(n-1)/n	S
15	$(2n-1)/n^2$	1/n	$(n-1)/n^2$	(S)

Occasion Contrast $(P[k] - P[\sim k])$:

Problem	Relative frequent	cy of target event for			
	Target occasion (k) (front plane)	Other occasions (~k) (other planes)	Contrast	Predicted attribution	
1	1 <i>/n</i>	1/n	0		
2	1/n	1	-(n-1)/n	ō	
3	$[(n-1)^2 + n]/n^2$	$[(n-1)^2 + n]/n^2$	0		
12	1	1/n	(n-1)/n	0	
15	1	$1/n^2$	$(n^2 - 1)/n^2$	0	

Appendix B (continued)

Person × Sti Problem	imulus Contrast {(P[i, j] - P[i, ~j]) - (P[~i, j] - P[~i, ~j])}: Relative frequency of target event for							
	[i, j]	[i, ~j]	[~i, j]	[~i, ~j]	[i, j]− [i, ~j]	[~i, j]- [~i, ~j]	Contrast	Predicted attribution
1	1	1	0	0	0	0	0	_
2	1	1	(n - 1)/n	(n - 1)/n	0	0	0	_
3	1	1	0	1	0	-1	1	$\mathbf{P} \times \mathbf{S}$
12	1	1/n	1	1/n	(n-1)/n	(n - 1)/n	0	
15	1	1/n	1/n	1/n	(n-1)/n	0	(n-1)/n	$\mathbf{P} \times \mathbf{S}$

$Person \times Occasion \ Contrast \ \{(P[i,k] - P[i, \sim k]) - (P[\sim i,k] - P[\sim i, \sim k])\}:$

Problem	Relative frequency of target event for							
	[i, k]	[i, ~k]	[~i, k]	[~i, ~k]	[i, k]− [i, ~k]	[∼i, k]− [∼i, ∼k]	Contrast	Predicted attribution
1	1	1	0	0	0	0	0	_
2	1	1	0	1	0	-1	1	P×O
3	1	1	(n-1)/n	(n-1)/n	0	0	0	
12	1	1/n	1	1/n	(n-1)/n	(n-1)/n	0	
15	1	1/n	1	0	(n-1)/n	1	-1/n	$\overline{(\mathbf{P} \times \mathbf{O})}$

$Stimulus \times Occasion \ Contrast \ \{(P[j,k]-P[j,\sim k]) - (P[\sim j,k]-P[\sim j,\sim k])\}:$

Problem								
	[j, k]	[j,~k]	[~j, k]	[~j, ~k]	[j, k]– [j, ~k]	[∼j, k]− [∼j, ∼k]	Contract	Predicted attribution
1	1/n	1/n	1/n	1/ <i>n</i>	0	0	0	_
2	1/n	1	1/n	1	-(n-1)/n	-(n-1)/n	0	_
3	1/n	1/n	1	1	0	0	0	
12	1	1	1	0	0	1	-1	$\overline{\mathbf{S} \times \mathbf{O}}$
15	1	1/n	1	0	(n - 1)/n	1	-1/n	$(\overline{\mathbf{S} \times \mathbf{O}})$

Received July 21, 1989

Revision received November 14, 1989

Accepted November 14, 1989