

A dual-process model of belief and evidence interactions in causal reasoning

JONATHAN A. FUGELSANG

Dartmouth College, Hanover, New Hampshire

and

VALERIE A. THOMPSON

University of Saskatchewan, Saskatoon, Saskatchewan, Canada

In three experiments, we examined how reasoners' preexisting beliefs about causal relations constrained their evaluation of covariation-based empirical evidence. Reasoners were presented with causal candidates that were a priori rated to be either believable or unbelievable, as well as information regarding the degree to which the cause and the effect covaried. Several findings supported the conclusion that preexisting beliefs about causal relations reflect knowledge of both causal mechanisms and covariation relations, that these sources of knowledge are represented independently and contribute independently to causal judgments, and that the evaluation of new empirical evidence is influenced differently by mechanism-based and covariation-based beliefs. Finally, we observed that reasoners were relatively accurate in evaluating the degree to which their judgments were sensitive to empirical evidence but were less able to judge how much their judgments were influenced by their prior beliefs. We present a dual-process model that provides a descriptive account of the boundary conditions for belief and evidence interactions in causal reasoning.

Reasoners' judgments of the causal relation between a putative causal candidate and an observed effect may be influenced by two variables: (1) the degree to which the candidate and the effect are observed to covary (e.g., Cheng & Novick, 1990) and (2) the reasoners' prior beliefs or knowledge about the two events in question (e.g., White, 1989). We investigated the sources of information that underlie causal beliefs and how these beliefs constrain the evaluation of covariation-based evidence.

A Brief Note on Terminology

In this article, *causal beliefs* will be defined as one's preexisting knowledge about the propensity of a given causal candidate to produce a given effect. We propose that these beliefs may be derived from several sources of information. We were specifically interested in two potential sources of causal beliefs: (1) long-term memory representations of covariation-based information and (2) long-term memory representations of mechanism-based properties of the relevant variables. We will refer to causal beliefs derived from covariation-based cues as *covariation-based beliefs* (or beliefs reflecting *covariation-based information*), whereas causal beliefs derived from mechanism-based cues will be referred to as *mechanism-*

based beliefs (or beliefs reflecting *mechanism-based information*). This variable will be referred to as *belief modality*.

In addition, we investigated how these beliefs influence the evaluation of new empirical evidence. In this series of experiments, empirical evidence was provided in the form of summary data that described the degree to which the two variables were observed to covary with each other in a hypothetical study. The empirical evidence provided to participants will be generally referred to as *covariation-based data* (or more descriptively, as the ΔP_c *manipulation*).

Covariation-Based Models

Some researchers claim that humans are "intuitive statisticians" (Peterson & Beach, 1967) who make causality judgments by using a normative strategy wherein the degree of covariation between a putative cause and its observed effect is computed. This model is based on the assumption that an event that exhibits a regularity of association with an effect (i.e., covaries with that effect) is more likely to be identified as a cause of that effect than is an event that does not exhibit a regularity of association. For example, eating peanuts is likely to covary more strongly with allergic reactions than is eating broccoli and, as such, will be more likely to be perceived as causal. Such covariation-based models of causation are the product of the Humean tradition of radical empiricism (Hume, 1739/1978), which posits that humans and other animals rely primarily on observable empirical cues to understand and explain causal sequences. Cur-

This research was supported by a postgraduate scholarship from the Natural Sciences and Engineering Research Council of Canada (NSERC) awarded to J.A.F. and by an NSERC discovery grant awarded to V.A.T. Correspondence should be addressed to J. A. Fugelsang, Department of Psychological and Brain Sciences, 6207 Moore Hall, Dartmouth College, Hanover, NH 03755 (e-mail: jonathan.a.fugelsang@dartmouth.edu).

rent versions of this view (e.g., Cheng, 1997) propose that causal attributions are based on the perceived degree of covariation between a candidate cause and an effect, as well as on the temporal and spatial relations between them.

Measuring the covariation between cause and effect. Summary information about the presence and absence of a putative cause and its effect is often represented in the form of a 2×2 contingency table. Figure 1 depicts the cells representing the four possible event conjunctions needed to compute an estimate of covariation. Covariation-based models typically assume that causal inferences are made by combining information from these four cells and deriving an estimate of the strength of the causal relation (e.g., Anderson & Sheu, 1995; Cheng, 1997; Cheng & Novick, 1990, 1992; Einhorn & Hogarth, 1986; Jenkins & Ward, 1965; Rescorla, 1968; Salmon, 1965; White, 2000, 2002b).

Proponents of this and other covariation-based models claim that the covariation between the cause and the effect is used as part of a causal model or schema to obtain a measure of causal strength between a putative cause and its effect (Cheng, 1997; Cheng & Novick, 1990, 1992). Take, for instance, the probabilistic contrast model (PCM; Cheng & Novick, 1990). This model and its successor, the power pc theory (Cheng, 1997), assume that causal judgments are based on two pieces of information: the probability of the effect's occurring when the cause is present [$P(e/c)$] and the probability of the effect's occurring when the cause is absent [$P(e/\sim c)$]. As is illustrated in Figure 1, these probabilities can be computed from the event conjunctions represented in a 2×2 contingency

table. A measure of covariation (ΔP_c) can then be computed, using the unidirectional contingency rule:

$$\Delta P_c = P(e/c) - P(e/\sim c). \tag{1}$$

If ΔP_c is positive, the candidate cause (c) should be judged to be a facilitatory factor in producing the effect (e). If ΔP_c is negative, the potential cause (c) should be judged to be an inhibitory factor. Finally, when ΔP_c is zero, the candidate cause (c) should be judged as a non-causal factor with respect to the observed effect (e).

Evidence for normative covariation-based models. The ΔP_c index is regarded by virtually all researchers in the field as a normatively appropriate index of the contingency between two binary variables (Kao & Wasserman, 1993). Furthermore, it has been proposed that the contingency-based model of causation appears to capture (at least in part) our everyday notion of how a cause and effect sequence might behave (Spellman, 1996). For example, the probability of getting lung cancer is greater if one smokes than if one does not smoke (i.e., ΔP_c is positive); this is often interpreted as evidence in favor of a causal relation between smoking and lung cancer, even though smoking is neither necessary nor sufficient for getting lung cancer. In this way, the fact that the two events covary supports a causal attribution, despite the fact that covariation does not guarantee that there is a causal link between them.

Laboratory evidence also suggests that reasoners make causal judgments that conform, at least approximately, to the ΔP_c contingency rule. Specifically, when reasoners are asked to judge the probability that a given cause produced a given effect, their likelihood judgment




Figure 1. A 2×2 contingency table representing the four possible event conjunctions required to calculate the conditional relation between an event and an outcome. The cells a, b, c, and d are the frequency of each conjunction and are used to calculate the conditional probability of the effect's (e) occurring in the presence of the cause (c) [$P(e/c)$] and its absence [$P(e/\sim c)$]. Note that ΔP_c can be expressed in terms of both a probability and an event frequency.

ments generally increase with increasing values of ΔP_c . This is the case regardless of whether the information needed to compute ΔP_c is presented in a 2×2 contingency table (Allan & Jenkins, 1980), is presented in a free-operant paradigm in which participants observe the putative causes and subsequent effects continuously in time (Lober & Shanks, 2000; Wasserman, Chatlosh, & Neunaber, 1983), or is summarized in sentences (Cheng & Novick, 1990).

Limitations of covariation-based models. Although there is much evidence to suggest that reasoners are sensitive to ΔP_c information, covariation-based models face two problems. The first is a problem in principle, in that covariation does not necessarily imply causation. Take, for example, the regular succession of day and night, where day is perfectly contiguous with night ($\Delta P_c = 1$) and appears temporally prior to night. Despite this perfect covariation, reasoners know that these two events are not causally linked. Although sequences such as this exhibit observable statistical characteristics similar to those for their valid causal counterparts, they lack the critical connection (i.e., causal mechanism) implied by a truly causal relation (Cheng, 1997).

Second, there are a number of research findings that challenge traditional contingency-based models. For example, several studies have demonstrated that participants will confidently infer a strong causal link between two events after observing only a single positive instance of a cause and an effect's co-occurring (Beasley, 1968; Boyle, 1960; Michotte, 1963). These reasoners cannot be using a derivative of the ΔP_c rule to inform their judgments, because ΔP_c requires at least two observations, one of which must include information about the probability of the effect's occurring in the absence of the cause [i.e., $P(e/\sim c)$]. Moreover, although ΔP_c assumes equal weighting for all four cells of a 2×2 contingency table, participants appear to weight Cells A and B more heavily than Cells C and D (Downing, Sternberg, & Ross, 1985; Schustack & Sternberg, 1981; Wasserman, Dornier, & Kao, 1990; White, 2002a). In addition, it appears that the order in which participants receive either confirmatory (Cells A and D) or disconfirmatory (Cells B and C) information can have a profound impact on participants' final judgments of causality (Collins & Shanks, 2002; Dennis & Ahn, 2001; Lopez, Shanks, Almaraz, & Fernandez, 1998) even though the order of presentation does not affect the objectively computed value of ΔP_c . As a result, there has been a move to modify the ΔP_c contingency rule (e.g., White, 2000, 2002b) to model human behavior more accurately.

Clearly, these findings indicate that participants do not rely on covariation as their only cue to causality. In their now classic review of causality judgments, Einhorn and Hogarth (1986) discussed a number of cues to causality, such as temporal order (Siegler & Liebert, 1974; Tversky & Kahneman, 1980), contiguity in time and space (Bullock, Gelman, & Baillargeon, 1982; Michotte, 1963), and similarity between cause and effect (Shultz & Ravinsky,

1977; Tversky, 1977), all of which may be used to support a causal attribution even when they conflict with covariation-based cues. In this view, covariation-based evidence is one of several cues to causality, rather than the primary or dominant cue (White, 1992).

Concept-Based Models

Whereas covariation-based models suggest that reasoners use empirical observation to evaluate causal hypotheses, concept-based models give priority to the role of acquired knowledge in evaluating causal relations. The lineage of such models can be traced back to Kant's (1781/1965) model of generative transmission. This view posits that causes not only covary with effects, but actually produce those effects. The term *generative transmission* refers to the transmission of energy from the cause to the effect such that the cause, through the nature of its properties, acts on the object, resulting in a causal outcome (Harre & Madden, 1975; Madden & Humber, 1974; White, 1989, 1995).

For example, the notion of generative transmission forms the core of the causal powers theory that was first proposed in philosophy by Harre and Madden (1975) and later applied to psychologically relevant questions by White (1989). Harre and Madden posited that causal powers are stable properties of objects, whose power to produce an effect is based on the "chemical, physical, or genetic natures of the entities involved" (p. 5); this power produces the effect only when the appropriate enabling conditions are present (White, 1989, 1995). For example, although peanuts may possess the causal power to produce allergic reactions, they will do so only if they are ingested. In this view, causal roles are defined conceptually, rather than through empirical associations. The assessment of causal hypotheses, therefore, is thought to be mainly a matter of (1) seeking some object believed to possess the power to produce the effect in question and (2) determining whether the appropriate enabling conditions are present to permit the power of the object to exert the effect.

Evidence for concept-based models. A number of studies have supported the hypothesis that reasoners' causal attributions reflect knowledge of causal powers. For example, when asked to evaluate the utility of various types of cues to test a causal hypothesis (e.g., covariation or temporal and spatial contiguity), reasoners reliably indicate that information regarding the causal mechanism will be the most informative (Shultz, 1982; Shultz, Fisher, Pratt, & Rulf, 1986). Moreover, when provided with the opportunity, reasoners do not spontaneously seek out covariation information between potential causal candidates and effects. Rather, individuals prefer to gather further information regarding the specific target events in question, to test hypotheses about possible underlying mechanisms (Ahn, Kalish, Medin, & Gelman, 1995; White, 1989). Taken together, these findings suggest that knowledge about a causal mechanism plays an important role in testing causal hypotheses and may even take priority over covariation-based data.

Combining Covariation-Based Data With Causal Beliefs

In summary, there is ample evidence that reasoners draw on at least two sources of information when making causal attributions. Reasoners are clearly sensitive to covariation-based data available in the environment; their judgments also reflect a need to understand the underlying causal mechanism that mediates the relation between cause and effect. However, despite the fact that both sources of knowledge clearly inform causal judgments, the tendency of most research to date has been to examine one or the other factor in isolation, or to pit covariation-based data and causal beliefs against each other as competing explanations for people's causal attributions. In our view, this approach has limited our understanding of causal judgments in two ways.

First, there is abundant evidence from other reasoning domains that shows that reasoners evaluate the information provided to them in light of their own beliefs and experiences and that their reasoning judgments are as sensitive to their prior beliefs as to the logical demands of the task (e.g., Evans, Handley, & Harper, 2001; Klauer, Musch, & Naumer, 2000; Newstead, Pollard, Evans, & Allen, 1992; Oakhill & Garnham, 1993; Thompson, Striemer, Reikoff, Gunter, & Campbell, 2003). Thus, it seems reasonable to expect that reasoners do not evaluate covariation-based data in an a theoretical manner but, instead, do so in light of their preexisting causal beliefs (Fugelsang & Thompson, 2000, 2001; White, 1995).

Second, the tendency to assume that either covariation- or mechanism-based cues have epistemological priority leads one to overlook the fact that these elements are closely related. Specifically, if I know that Cause A has the causal power to produce Effect B, I will almost certainly believe that A and B covary, even if I have not had the opportunity to observe such covariation empirically. For example, on the basis of my causal beliefs regarding the sharpness of tacks and the delicacy of the stomach lining, I am prepared to believe that ingesting tacks will covary with severe abdominal pain, even though I have never had the opportunity to observe the co-occurrence of these events. Conversely, knowing that two events covary can lead one to assume that there is a causal mechanism linking the two. For example, a great deal of research was carried out to discover a causal mechanism linking lung cancer to smoking on the basis of early findings that established that these events were correlated. In summary, causal beliefs likely reflect both covariation-based information and mechanism-based information, regardless of whether causal beliefs are initially formed by observing cues to covariation (e.g., Cheng & Lien, 1995; Young, 1995) or by inferring causal mechanisms (e.g., White, 1989).

Evidence for cue combination. To date, there have been few studies that have examined how reasoners combine covariation-based data with their causal beliefs. Much of the existing research concerns *when* covariation entails causality (e.g., Lien & Cheng, 2000; Waldmann, 1996;

Waldmann & Hagmayer, 2001) or *how* causal beliefs can moderate or downplay the importance of covariation-based data (Ahn et al., 1995; Garcia, McGowan, Ervin, & Koelling, 1968; Michotte, 1963; White, 1989). The data that are available suggest that these sources of information contribute interactively, as opposed to additively, to causal judgments (Fugelsang & Thompson, 2000). For example, participants in Fugelsang and Thompson's (2000) experiments judged the likelihood that a candidate cause produced a given effect following orthogonal manipulations of belief level and ΔP_c . We observed that the belief level and ΔP_c manipulations interacted, so that the effect of ΔP_c was larger for believable than for unbelievable candidates. We interpreted these findings to be consistent with White's (1989) hypothesis that mechanism-based beliefs restrict the set of candidates about which covariation-based data are sought out and considered.

We report three experiments designed to extend these findings. Our first goal was to test the hypothesis that mechanism-based beliefs per se are used to restrict the set of causal candidates about which covariation-based data are considered. In addition, we wished to test a second assumption of White's (1989) causal power theory. White (1989) proposed that mechanism-based beliefs are applied in an unconscious heuristic manner, rather than in a strategic thoughtful manner. Specifically, he assumed that mechanism-based beliefs are recruited automatically and that their influence on subsequent causal attributions are thus beyond one's conscious control. Evans and Over (1996, 1997) have made similar arguments in the case of deductive reasoning. They proposed that the recruitment of beliefs to problems of deduction may occur unconsciously and, thus, may be beyond the realm of introspection. In this view, strategic analytic processes are proposed to operate after heuristic belief-based processes have identified elements in the environment for subsequent processing. We therefore predicted that reasoners would be unable to introspect accurately on the degree to which their causal judgments were affected by their prior mechanism-based beliefs. In contrast, we hypothesized that covariation-based data would be used strategically and analytically, after belief-based processes have been applied, and thus would be available for introspection.

EXPERIMENT 1

Our goal in the present experiments was to vary mechanism-based beliefs independently of covariation-based beliefs. In Experiment 1, we did this by presenting participants with unknown candidates (i.e., *Substance X*) and manipulating (1) the degree to which they would be believed to covary with an effect and (2) the degree to which they possessed a causal mechanism capable of producing the given effect. In addition, we presented the reasoners with information needed to compute ΔP_c and examined how the use of these covariation-based data

changed across different levels of belief modality. To test our hypothesis about the strategic use of belief-based information and covariation-based data, we asked the participants to record how much they had relied on the believability of the causal candidates, as well as on the covariation-based data provided to them (see Jack & Roepstorff, 2002, for a recent analysis of similar retrospective introspection techniques).

Method

Participants. One hundred seventy-one first-year psychology students from the University of Saskatchewan, with a mean age of 20.2 years (range, 18–44), completed the study in partial fulfillment of a course requirement.

Design. This was a $2 \times 2 \times 2$ mixed design with belief level (low vs. high) and covariation-based data ($\Delta P_c = .1$ vs. $.9$) as within-subjects variables; belief modality (mechanism based vs. covariation based) was a between-subjects variable.

Materials and Procedure. For this experiment, we used four story scenarios, each presented on a separate page in a small booklet. Each story scenario concerned a different effect in need of explanation: fatigue, slippery roads, flowers blooming, and lung cancer (a complete set of stimuli is presented in the Appendix). Each scenario described causal links that were either believable or unbelievable. Believability ratings were pretested with a separate sample of 91 participants. These ratings, along with participants' estimates of the degree to which the putative causes and effects covaried in the natural environment, appear in Table 1. The critical manipulation in this experiment was the manipulation of the source of those beliefs (i.e., the belief modality). Half of the participants received information describing how the cause and the effect had covaried in the past (i.e., *covariation-based beliefs*); the remaining half received information describing a causal mechanism (i.e., *mechanism-based beliefs*). Each participant solved four problems, one in each belief level $\times \Delta P_c$ cell. The four story scenarios and the two ΔP_c conditions were fully counterbalanced, so that each scenario appeared equally often under each covariation contingency. Each problem began with a brief introductory paragraph that explained an event that had happened and a possible cause for that event. The possible cause was an unknown candidate (e.g., *Substance X*) that was arbitrarily linked to the effect. For example, in the *slippery roads* scenario, the participants were provided with the following paragraph:

Imagine you are a municipal researcher who works for a local township. Recently, you have discovered that a small section of town is showing a dramatic increase in the number of slippery roads. You have a hypothesis that the slippery roads may be due to the recent use of "substance X" in the mixture of asphalt.

The belief modality information was then provided. Specifically, the nature of *Substance X* was then described either in terms of the

inherent properties of the substance (mechanism based) or in terms of its statistical relation with the effect (covariation based). For example, in the mechanism-based condition, the participants were told the following: "You discover that 'substance X' is a chemical additive that alters the freezing point of precipitation such that rain freezes at a warmer temperature." This information provides the reasoner with information regarding the causal power of *Substance X* to elicit the observed effect without providing any information about past statistical regularities. The participants in the covariation-based condition were told that "past research has shown that 'substance X' and slippery roads are strongly correlated, although the reasons for this relation are not fully understood." This information provides the reasoner with information regarding past statistical regularities without revealing the exact nature of the causal candidate.

At this point, the participants made an initial likelihood judgment. They were asked to judge the likelihood that the putative causal candidate was responsible for the observed effect, using an 11-point Likert scale that ranged from 0 (*definitely not likely*) to 10 (*definitely likely*), with 5 (*somewhat likely*) as the midpoint. Note that this initial likelihood judgment is based solely on knowledge of the believability of the candidate cause in the absence of covariation-based data. We shall refer to these judgments as baseline *perceived causal efficacy* (PCE).

Immediately following the baseline PCE judgment, we presented the covariation-based data needed to compute ΔP_c . This information was presented in a frequency format. For example, in the highly contingent ($\Delta P_c = .9$) *slippery roads* scenario, the participants were provided with the following new empirical evidence:

In order to test this theory, you investigate 10 townships that had "substance X" and 10 townships that did not have "substance X." A thorough investigation revealed the following information: of the 10 townships that had "substance X," 9 experienced slippery roads; of the 10 townships that did not have "substance X," 0 had slippery roads.

In the high-contingency condition, ΔP_c was set equal to $.9$; in the low-contingency condition, ΔP_c was equal to $.1$. The marginal totals (i.e., total number of observations for which the cause was present or absent) were set at 10 for all levels of ΔP_c . The value of ΔP_c was manipulated by changing the probability of the effect's occurring in the presence of the cause [i.e., $P(e/c)$], as is illustrated in Table 2.

This information was followed by another Likert scale, which was identical to the first one. The participants were then asked to make a second PCE judgment reflecting all sources of information given (belief based and evidence based). Finally, for each scenario, the participants were asked to indicate (1) how heavily they had weighted their causal beliefs and (2) how heavily they had weighted the new covariation-based data. These judgments were also made using an 11-point Likert scale that ranged from 0 (*not at all*) to 10 (*extremely heavily*) and will be referred to as belief-weighting and ΔP_c -weighting, respectively.

All the participants were tested in a single session. The participants were randomly assigned to one of the two belief modality

Table 1
Mean Plausibility Ratings and Covariation Estimates [i.e., both $P(e/c)$ and $P(e/-c)$]
for the Scenarios Used in Experiments 1 and 2

Scenario	Candidate Cause	Plausibility	$P(e/c)$	$P(e/-c)$
Cancer	taking iron supplements	2.57	.17	.20
Exam success	paper writing success	5.55	.73	.38
Fatigue	insomnia	7.96	.85	.22
Fever	having chills	3.60	.66	.17
Flowers blooming	red pots	0.23	.75	.72
Slippery roads	ice storms	9.28	.96	.19

Note—All six items were used in Experiment 2; only those appearing in boldface were used in Experiment 1. These data were gathered in an independent sample of 91 participants.

Table 2
Event Frequencies for the Computation of the ΔP_c Values Used in Experiments 1, 2, and 3

Frequencies				ΔP_c Computations			
<i>ce</i>	<i>c-e</i>	<i>~ce</i>	<i>~c~e</i>	$P(e/c)$	–	$P(e/~c)$	ΔP_c
9	1	0	10	9/10	–	0/10	.9
1	9	0	10	1/10	–	0/10	.1

Note—*c*, cause; *e*, effect; *ce* represents the number of times that the cause and the effect co-occurred; *c-e* represents the number of times that the cause occurred in the absence of the effect; *~ce* represents the number of times that the effect occurred in the absence of the cause; *~c~e* represents the number of times that the effect was absent when the cause was absent.

conditions. The materials for this experiment were included in a set of unrelated experiments on other topics. Instructions were all written, so the experimenter gave only a brief introduction informing the participants that they should complete the questions in the order in which they appeared and that they should work at their own pace. The participants were encouraged to ask any questions they might have at any time during the experiment.

Results and Discussion

The results will be presented in three sections. The first section will deal with the baseline PCE judgments (i.e., the judgments made prior to the provision of the covariation-based data). In the second section, the final PCE judgments will be analyzed as a function of belief level, belief modality, and ΔP_c . In the final section, we will compare the participants' perceived use of causal cues to their actual use of causal cues by comparing their subjective belief- and ΔP_c -weighting judgments with their actual judgments. The α level for all the statistical tests was set at .05 (two-tailed) unless otherwise stated. Effect size estimates were computed using partial η^2 .

Baseline PCE judgments. The first series of analyses served as a manipulation check to verify that the

mechanism-based and covariation-based belief manipulations affected PCE judgments in the predicted manner. Recall that these baseline PCE judgments were made after the participants had been provided with information regarding a causal mechanism or covariation-based information but prior to the presentation of new empirical evidence (i.e., the ΔP_c manipulation). The baseline PCE judgments should, therefore, vary across levels of belief, but not across levels of ΔP_c . Figure 2 presents the mean PCE judgments (baseline and final) for the two belief level conditions (low belief and high belief) and the two ΔP_c conditions (.1 and .9) as a function of belief modality (mechanism and covariation based).

The baseline PCE judgments were analyzed using a $2 \times 2 \times 2$ (belief level \times $\Delta P_c \times$ belief modality) mixed analysis of variance (ANOVA). There was an expected main effect of belief level [$F(1,169) = 385.56$, $MS_e = 4.01$, $\eta^2 = .68$] such that the participants' baseline PCE judgments were higher when the candidates were believable ($M = 7.2$) than when they were not ($M = 4.2$). This main effect was qualified, however, by an interaction between belief level and belief modality [$F(1,169) = 10.03$, $MS_e = 4.01$, $\eta^2 = .02$]: The difference between high belief and low belief was larger for the covariation-based ($M = 3.5$) than for the mechanism-based ($M = 2.5$) condition. The effect of ΔP_c ($F = 3.68$), the $\Delta P_c \times$ belief modality interaction ($F = 2.29$), and the three-way interaction ($F < 1$) were not significant (all $ps > .05$). Therefore, the belief manipulation successfully influenced the participants' belief judgments and was not contaminated by the later ΔP_c manipulation.

Final PCE judgments. The final PCE judgments were analyzed using a $2 \times 2 \times 2$ (belief level \times $\Delta P_c \times$ belief modality) mixed ANOVA. As was expected, PCE judgments were higher for believable ($M = 6.1$) than for

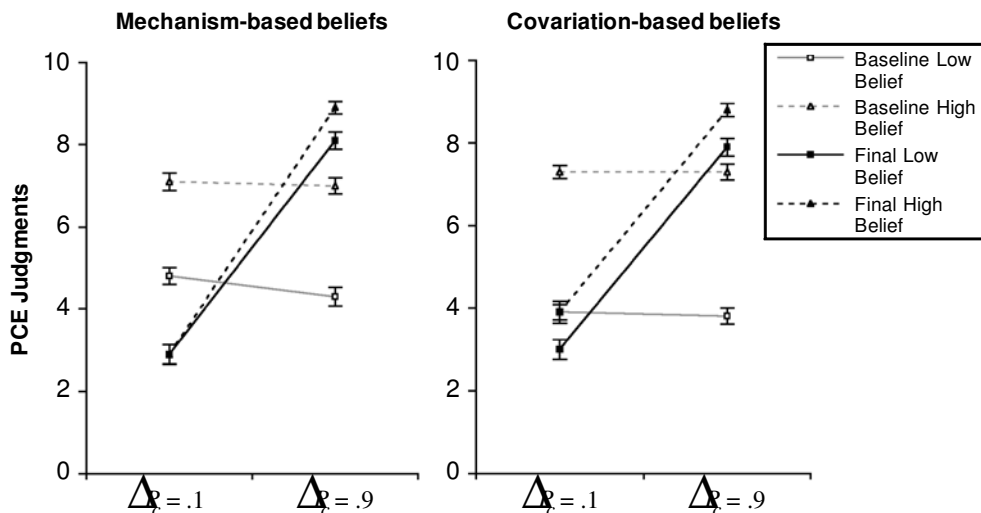


Figure 2. Experiment 1: mean perceived causal efficacy (PCE) judgments for the two belief level conditions (low belief and high belief) and the two ΔP_c conditions (.1 and .9) as a function of belief modality (mechanism vs. covariation based).

unbelievable ($M = 5.5$) candidates [$F(1,169) = 25.45$, $MS_e = 2.52$, $\eta^2 = .13$]; PCE judgments were also higher when $\Delta P_c = .9$ ($M = 8.4$) than when $\Delta P_c = .1$ [$M = 3.2$; $F(1,169) = 771.28$, $MS_e = 6.06$, $\eta^2 = .82$]. As with the baseline PCE judgments, there was also a belief level \times belief modality interaction [$F(1,169) = 4.26$, $MS_e = 2.52$, $\eta^2 = .03$]. The difference between believable and unbelievable candidates was larger for the covariation-based ($M = 0.9$) than for the mechanism-based ($M = 0.4$) belief modality condition. The main effect of belief modality ($F = 1.15$), the $\Delta P_c \times$ belief level interaction ($F = 2.36$), and the $\Delta P_c \times$ belief modality interaction ($F = 3.28$) were not significant (all $ps > .05$).

Importantly, the predicted interaction between belief level, ΔP_c , and belief modality [$F(1,169) = 3.74$, $MS_e = 1.86$, $p = .055$, $\eta^2 = .02$] was marginally reliable. To interpret this three-way interaction, separate belief level \times ΔP_c ANOVAs were computed for the two belief modality conditions. As was predicted, in the mechanism-based belief modality condition, the effect of ΔP_c was larger when the candidate was highly believable (M difference = 5.9) than when the candidate was of low belief [M difference = 5.2; $F(1,169) = 6.77$, $MS_e = 1.68$, $\eta^2 = .07$]. In contrast, however, in the covariation-based belief modality condition, belief level and ΔP_c contributed additively to PCE judgments ($F < 1$, $\eta^2 < .01$); the effect of ΔP_c was equivalent for believable (M difference = 4.9) and unbelievable (M difference = 4.9) candidates. This null effect occurred despite adequate power to detect an interaction of similar magnitude to that in the previous analysis [$1 - \beta$ for (M difference $\geq .72$) = .83, $\alpha = .05$, one-tailed].

These findings support the conclusion that mechanism-based and covariation-based beliefs are distinct and contribute in different ways to causality judgments (Fugelsang & Thompson, 2000). Specifically, when ΔP_c is evaluated in light of covariation-based beliefs, it appears that the

two sources of information are simply added together, perhaps because they are derived from the same modality. In contrast, when ΔP_c is evaluated in light of mechanism-based beliefs, it is weighted differently depending on the believability of the candidate. Specifically, covariation-based data are weighted more heavily for believable than for unbelievable causal candidates. This pattern is consistent with our hypothesis that mechanism-based beliefs are used to restrict the set of candidates for which covariation information is considered (Fugelsang & Thompson, 2000; White, 1989).

Actual/subjective calibration. In the final series of analyses, we examined the accuracy of the participants' reported use of belief-based information and covariation-based data. To do this, two scores were computed that represented (1) actual use of belief-based information and covariation-based data and (2) subjective use of belief-based information and covariation-based data. The actual use of belief-based information and covariation-based data was determined by creating difference scores that represented the degree to which the reasoners' judgments changed as a function of belief level and ΔP_c across problems (i.e., the difference in final PCE scores between the high- and the low-belief conditions and the difference in final PCE scores between the high- and low- ΔP_c conditions). Likewise, we computed two scores that represented the degree to which the reasoners believed that their judgments changed as a function of the belief level manipulations and the ΔP_c manipulations; this was done by averaging the participants' subjective weighting judgments of belief-based information (i.e., how much they thought that they had made use of the information regarding the believability of the candidate) and covariation-based data (i.e., how much they thought that they had made use of the new empirical evidence) across problems. These data are plotted in Figure 3.

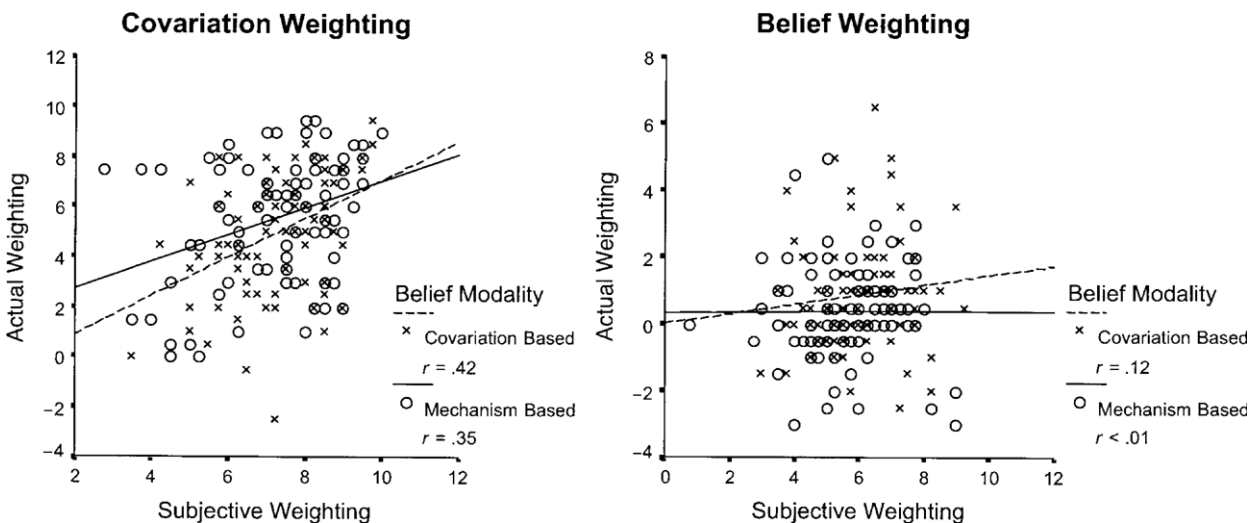


Figure 3. Experiment 1: individual calibration of participants' actual use of belief- and covariation-based cues with their subjective use of belief- and covariation-based cues.

These analyses revealed several interesting findings. First, the participants' subjective weighting of covariation-based data was positively associated with their actual use of covariation-based data when making their final PCE judgments [$r(169) = .38, p < .01$]. In other words, the participants were able to gauge their use of covariation-based data: The participants who said that their judgments were influenced by the covariation-based data did in fact change their judgments across problems as a function of covariation-based data. In contrast, however, the participants' subjective weighting of belief-based information was not correlated with their actual use of belief-based information when making their final PCE judgments [$r(169) = .07, p = .34$], despite adequate power to detect a correlation of moderate size [$1 - \beta$ for ($r \geq .3$) = .97]. It appears, therefore, that the participants were unable to gauge their use of belief-based cues and were unaware of the influence that this information had on their final PCE judgments. Performing a simple comparison between the slopes of the belief (actual vs. subjective) and covariation (actual vs. subjective) calibration estimates corroborated these analyses. This analysis revealed that these two slopes were reliably different from each other [$t(168) = 3.32, p < .05$] and, thus, represented dissociable processes. However, the slopes did not vary as a function of the belief modality manipulation (all t s < 1 , p s $> .05$), suggesting that the influence of causal beliefs may be unavailable for introspection, regardless of the source of those beliefs.

In summary, we report two novel findings concerning the relation between preexisting beliefs and the evaluation of new covariation-based data. Using unfamiliar causal candidates, we demonstrated that reasoners weight covariation-based data more heavily for believable than for unbelievable candidates, but only when there is a causal mechanism that links the cause to the effect; otherwise, prior beliefs and new data contribute additively to causal judgments. We also demonstrated that the reasoners are relatively accurate at evaluating how much their causal judgments reflect the impact of covariation-based data but seem to have little insight into how much their causal judgments reflect the impact of their prior beliefs. The purpose of the next two experiments was to extend and replicate each of these findings, using familiar causal relations.

EXPERIMENT 2

In Experiment 1, we sought to manipulate mechanism-based and covariation-based beliefs by using unknown candidates. In Experiment 2, in contrast, we used familiar causal candidates; these candidates varied in whether there was a mechanism that mediated the relation and in the degree to which the cause and the effect were believed to covary. A second goal of this experiment was to extend our analyses to judgments of covariation. Recent research has suggested that the inferential processes associated with judgments of covariation are different from those associated with judgments of causality (e.g., Mandel & Lehman, 1998; White, 2001). Consequently, we wanted

to determine whether the interaction we observed between ΔP_c and belief level would be observed when reasoners were asked to judge the degree of covariation between cause and effect, as opposed to judging the strength of the causal relation (i.e., PCE judgments).

Thus, the participants in Experiment 2 were provided with causal scenarios in which the causal candidates varied in terms of (1) the belief modality (covariation vs. mechanism based), (2) the belief level (high vs. low), and (3) the strength of covariation-based data ($\Delta P_c = .1$ vs. $.9$). In addition, the participants made one of two types of judgments. Specifically, the participants either (1) made PCE judgments or (2) estimated the degree to which the cause and the effect covaried (hereafter referred to as COV estimates). We hypothesized that the belief level $\times \Delta P_c$ interaction would occur only when beliefs were represented in terms of mechanism-based information and individuals were making PCE judgments, rather than COV estimates.

Method

Participants. One hundred thirty-four first-year psychology students from the University of Saskatchewan, with a mean age of 18.8 years (range, 17–49), completed the study in partial fulfillment of a course requirement.

Design. This was a 3 (belief level) \times 2 (covariation-based data) \times 2 (judgment type) mixed design, with belief-level (low vs. high) and covariation-based data ($\Delta P_c = .1$ vs. $.9$) as within-subjects variables and judgment type (PCE judgment vs. COV estimate) as a between-subjects variable.

Materials and Procedure. In this experiment, the participants were presented with six story scenarios, one on each page of a small booklet. Each scenario presented an event in need of explanation (i.e., fatigue, slippery roads, flowers blooming, lung cancer, fevers, and exam success) and a putative cause. The scenarios were selected on the basis of ratings obtained in the same pretesting session as that noted in Experiment 1. These ratings appear in Table 1. Because covariation-based beliefs and mechanism-based beliefs are highly correlated, we were unable to manipulate these dimensions orthogonally using naturalistic materials. Consequently, we presented causal candidates having one of three relations to the target event: The candidate was highly correlated with the effect but was not believed to have a causal mechanism to link it to the effect (e.g., having chills causing a fever), the candidate was correlated with the effect and was believed to possess a causal mechanism (e.g., insomnia causing fatigue), or the candidate was not correlated and was not believed to have a mechanism (e.g., taking iron supplements causing lung cancer). The latter condition served as the low-belief control condition.

Each putative cause was presented with covariation-based data that were either highly contingent with the effect ($\Delta P_c = .9$) or of low contingency with the effect ($\Delta P_c = .1$). The causal scenarios and the ΔP_c manipulations were fully counterbalanced so that each scenario appeared equally often under both covariation contingencies. As in Experiment 1, the marginal totals (i.e., total number of observations for which the cause was present or absent) were set at 10 for all levels of ΔP_c . The value of ΔP_c was manipulated by changing the probability of the effect's occurring in the presence of the cause [i.e., $P(e|c)$], as is illustrated in Table 2. Finally, the participants were asked to make judgments regarding either (1) how likely it was that the given candidate caused the given effect or (2) the degree to which the cause and the effect covaried.

An example of the scenarios given to the participants who made PCE judgments is as follows. The participants were first provided with a brief introductory paragraph that explained the event that had

happened and a possible cause for that event. For example, in the *fatigue scenario*, the participants were provided with the following paragraph:

Imagine you are a doctor trying to determine the cause of fatigue in a group of patients. You have a hypothesis that the fatigue may be due to insomnia. In order to test this theory, you investigate 10 patients that had insomnia and 10 patients that did not have insomnia. A thorough investigation revealed the following information: of the 10 patients that had insomnia, 9 were fatigued; of the 10 patients that did not have insomnia, 0 were fatigued.

At this point, the participants judged the likelihood that the putative candidate was causally responsible for the observed effect, using an 11-point Likert scale that ranged from 0 (*definitely not likely*) to 10 (*definitely likely*), with 5 (*somewhat likely*) as the midpoint.

The participants who made COV estimates were provided with the identical scenarios, save that references to the evaluation of causal relation changed. The critical changes are shown in *italics* in the example. For example, in the *fatigue scenario*, the participants were provided with the following paragraph:

Imagine you are a doctor trying to determine if *there is a relationship between insomnia and fatigue in a group of patients*. In order to test this theory, you investigate 10 patients that had insomnia and 10 patients that did not have insomnia. A thorough investigation revealed the following information: of the 10 patients that had insomnia, 9 were fatigued; of the 10 patients that did not have insomnia, 0 were fatigued.

The participants then judged the strength of the relation between the two variables, using a similar 11-point Likert scale that ranged from 0 (*not related*) to 10 (*very strong relationship*), with 5 (*moderate relationship*) as the midpoint. The procedure was otherwise identical to that in Experiment 1.

Results and Discussion

Figure 4 presents the mean PCE judgments and COV estimates for the two belief level conditions (low belief and high belief) and the two ΔP_c conditions (.1 and .9) as a function of belief modality (mechanism based vs. covariation based). In the first set of analyses, we examined the effects of belief level, ΔP_c , and belief modality. The PCE judgments were analyzed using two separate

2×2 (belief level $\times \Delta P_c$) ANOVAs. In the first 2×2 ANOVA, scenarios describing candidates having a believable causal mechanism were contrasted to the low-belief control. In the second 2×2 ANOVA, scenarios describing candidates that were correlated with the effect but that lacked a believable causal mechanism were contrasted to the low-belief control. Parallel analyses were conducted for the COV estimates.

To simplify the description of these analyses, the four sets of analyses will be summarized together in terms of (1) main effects of belief level, (2) main effects of ΔP_c , and (3) the presence or absence of a belief level $\times \Delta P_c$ interaction. First, main effects of belief level were found for both belief modality manipulations and both judgment types (all $F_s > 5.6$, $\eta^2_s > .08$). That is, under all the conditions, the participants' judgments (PCE and COV) were higher when the candidate cause was highly believable ($M = 5.7$) than when it was of low believability ($M = 4.5$), regardless of how those beliefs were manipulated. In addition, main effects of ΔP_c were found for both belief modality manipulations and both judgment types (all $F_s > 234.68$, $\eta^2_s > .79$). That is, under all the conditions, the participants' judgments (PCE and COV) were higher when the candidate was highly contingent ($M = 8.2$) than when it was weakly contingent ($M = 2.1$), regardless of how those beliefs were manipulated.

Importantly, however, the belief level $\times \Delta P_c$ interaction was significant only when beliefs were manipulated in terms of mechanism-based information and the participants were making PCE judgments [$F(1,63) = 6.20$, $MS_e = 2.83$, $\eta^2 = .09$]. The form of this interaction was as expected: The effects of the ΔP_c manipulation were larger when the candidate was highly believable (M difference = 6.3) than when the candidate was of low believability (M difference = 5.3). Under all other conditions, the test for the belief level $\times \Delta P_c$ interaction was additive (largest $F =$

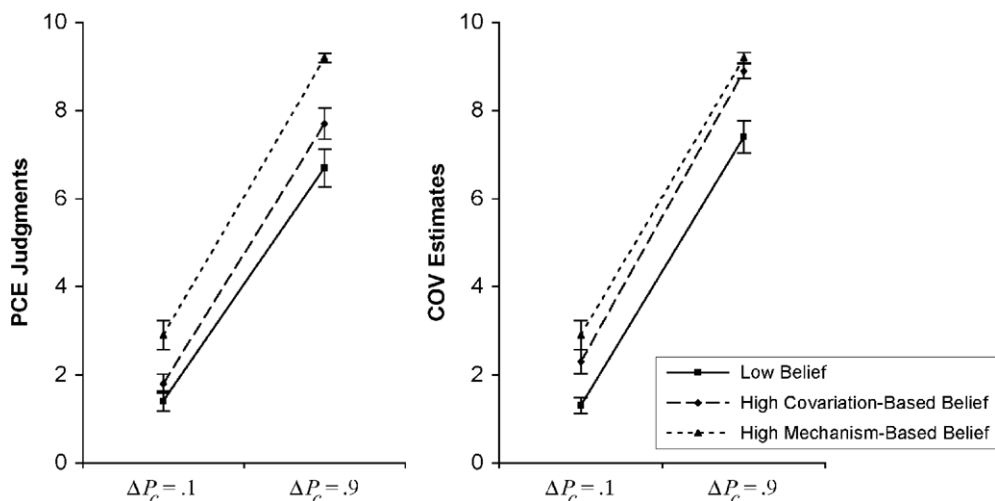


Figure 4. Experiment 2: mean perceived causal efficacy (PCE) judgments and covariation (COV) estimates for the two belief level conditions (low belief and high belief) and the two ΔP_c conditions (.1 and .9) as a function of belief modality (mechanism vs. covariation based).

1.2, $\eta^2 = .02$), despite adequate power to detect an interaction of magnitude similar to that observed for PCE judgments when beliefs were represented by mechanism-based information [$1 - \beta$ s for (M difference ≥ 1.05) = .83, $\alpha = .05$, one-tailed].

In summary, it appears that the interaction between belief level and ΔP_c is observed under restricted circumstances: (1) when judgments about the perceived causal efficacy of candidates (i.e., PCE judgments) are made and (2) when beliefs are represented in terms of mechanism-based information. On the other hand, if beliefs are represented in terms of long-term memory representations of covariation-based information, or when people are judging the degree to which two variables are correlated (i.e., COV estimates), the relation between belief level and ΔP_c is additive.

EXPERIMENT 3

The goal of Experiment 3 was to provide a replication of the introspection data reported in Experiment 1. In addition, we wanted to rule out a possible alternative interpretation of these findings. Recall that we found reasoners to be more accurate when introspecting about the contribution of covariation-based data to their causal judgments than when introspecting about the contribution of their prior beliefs. The interpretation of this finding is potentially compromised by the fact that the effect of covariation-based data on both actual and subjective judgments (M s = 5.2 and 7.2) was larger than the corresponding effect of beliefs [M s = 0.7 and 5.8; $t(170) = 19.25$, $SE = 0.24$, and $t(170) = 9.79$, $SE = 0.14$, respectively]. The restricted range for the belief-based measures may have artificially attenuated the correlation between the actual and the subjective measures. If this were the case, it would compromise our conclusions regarding the phenomenological characteristics of belief- and covariation-based processing.

Thus, a second goal of Experiment 3 was to test whether this dissociation still occurs when the effects of belief on both actual and subjective judgments are comparable in size to the effects of covariation-based data. Fugelsang and Thompson (2000, Experiment 2) observed that presenting the belief-based information prior to the covariation-based data reduced the size of the belief effect. Thus, one explanation for the relatively small belief effects observed in Experiment 1 was that the belief-based information was presented prior to the covariation-based data. Consequently, in the present experiment, we presented the belief-based information and covariation-based data simultaneously in order to reduce any possible recency effects associated with presenting the covariation-based data last.

Method

Participants. Two hundred forty-two first-year psychology students from the University of Saskatchewan, with a mean age of 19.9 years (range, 16–54), completed the study in partial fulfillment of a course requirement.

Design. This was a 2×2 design with belief level (low vs. high) and covariation-based data ($\Delta P_c = .1$ vs. $.9$) as within-subjects variables.

Materials and Procedure. Four story scenarios, adapted from Fugelsang and Thompson (2000), were presented on separate pages in a small booklet. The four scenarios concerned the following effects in need of explanation: depleted fish populations, car accidents, car start failures, and allergic reactions. As in Experiments 1 and 2, each scenario contained causal links that were either believable or unbelievable (see Fugelsang & Thompson, 2000, for pretested believability ratings, using a separate sample of 64 participants).

In addition, we accompanied each putative cause with covariation-based data that supported a relation that was either highly contingent ($\Delta P_c = .9$) or of low contingency ($\Delta P_c = .1$). As was the case in Experiment 2, the belief- and covariation-based information was presented simultaneously, and a single judgment (in this case, a PCE judgment) was made after both the belief-based information and the covariation-based data had been presented.

The four scenarios and two ΔP_c conditions were fully counter-balanced so that each scenario appeared equally often under both covariation contingencies. The marginal totals (i.e., total number of observations for which the cause was present or absent) were set at 10 for all levels of ΔP_c . As was the case in Experiments 1 and 2, the value of ΔP_c was manipulated by changing the probability of the effect's occurring in the presence of the cause [i.e., $P(e|c)$], as is illustrated in Table 2.

In addition to making PCE judgments, the participants were asked to indicate (1) how heavily they had weighted the belief-based information and (2) how heavily they had weighted the covariation-based data, using an 11-point Likert scale that ranged from 0 (*not at all*) to 10 (*extremely heavily*). The procedure was otherwise identical to that in Experiments 1 and 2.

Results and Discussion

The results will be presented in three separate sections. In the first section, we will examine the magnitude of the actual and subjective belief and ΔP_c effects. In the second section, we will analyze the PCE judgments as a function of the belief level and ΔP_c manipulations. The final section will present the calibration between the actual and the subjective uses of the causal cues.

The effects of belief level and ΔP_c . In the first set of analyses, we examined the degree to which we were successful in equalizing the effects of belief level and ΔP_c . These analyses revealed that the effects of belief level ($M = 3.7$) and of ΔP_c ($M = 3.8$) on actual judgments were comparable ($t < 1$), as were the subjective evaluations of the effect of belief ($M = 5.6$) and ΔP_c ($M = 5.6$; $t < 1$). Thus, we were successful in equating the effects of belief and ΔP_c on both actual and subjective judgments.

PCE judgments. The PCE judgments were analyzed using a 2×2 (belief level $\times \Delta P_c$) repeated measures ANOVA. Figure 5 presents the mean PCE judgments for the two belief level conditions (low belief and high belief) and the two ΔP_c conditions (.1 and .9). As was the case in Experiments 1 and 2, there was a main effect of belief level [$F(1,241) = 515.31$, $MS_e = 6.33$, $\eta^2 = .68$], such that PCE judgments were higher for believable ($M = 5.9$) than for unbelievable ($M = 2.3$) candidates. There was also an effect of ΔP_c [$F(1,241) = 536.25$, $MS_e = 6.67$, $\eta^2 = .69$], such that the candidates were given higher PCE judgments when they were highly contingent ($M = 6.0$) than when they were weakly contingent ($M = 2.2$). There

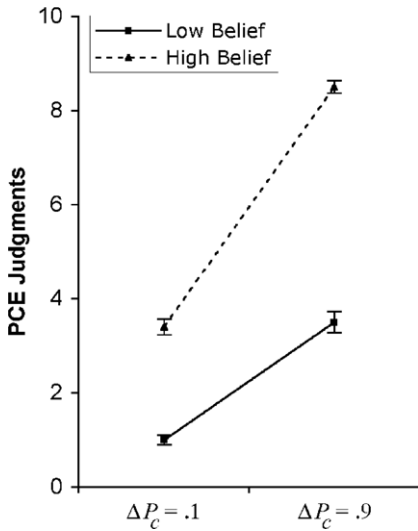


Figure 5. Experiment 3: mean perceived causal efficacy (PCE) judgments for the two belief level conditions (low belief and high belief) and the two ΔP_c conditions (.1 and .9).

was also a belief level $\times \Delta P_c$ interaction [$F(1,241) = 77.93$, $MS_e = 4.88$, $\eta^2 = .24$], such that the effects of ΔP_c were larger when the causal candidates were believable (M difference = 5.1) than when they were of low belief (M difference = 2.6).

Actual/subjective calibration. In the final series of analyses, we examined the extent to which the participants were aware of the influence of belief-based information and covariation-based data. These analyses paralleled those of Experiment 1. For each individual, scores representing the actual use of belief-based information were computed by subtracting the mean PCE judgment made in the high- and low-belief conditions;

comparable scores were computed for the two ΔP_c conditions. The subjective estimates were obtained by averaging the relevant ratings across problems. These data are plotted in Figure 6.

As was the case in Experiment 1, the participants' subjective weighting of covariation-based data was positively associated with their actual use of covariation-based data [$r(240) = .34$, $p < .01$]. In contrast, unlike Experiment 1, the subjective use of belief-based information was mildly associated with actual use of belief-based cues [$r(240) = .18$, $p < .01$]. Therefore, when baseline belief effects were substantially increased, the participants' awareness of their influence did increase slightly. Importantly, however, the two correlations differed reliably [$t(239) = 2.30$, $p < .05$]. Thus, the participants appeared to be less accurate when introspecting about the contribution of prior knowledge than when introspecting about the contribution of covariation-based data to their PCE judgments.

GENERAL DISCUSSION

Our findings support two conclusions. First, beliefs about causal relations reflect at least two independent nonredundant sources of information: knowledge about the degree to which the cause and the effect covary and knowledge about causal mechanisms that mediate between cause and effect. These two sources of knowledge appear to be represented independently and appear to contribute differently to the subsequent evaluation of covariation-based data. Specifically, we found that covariation-based data were weighted more heavily for believable than for unbelievable candidates, but only when those candidates were believed to have a causal mechanism linking them to their effect *and* when the reasoners were making causal judgments, as opposed to estimating covariation. These findings held for both familiar and unfamiliar causal re-

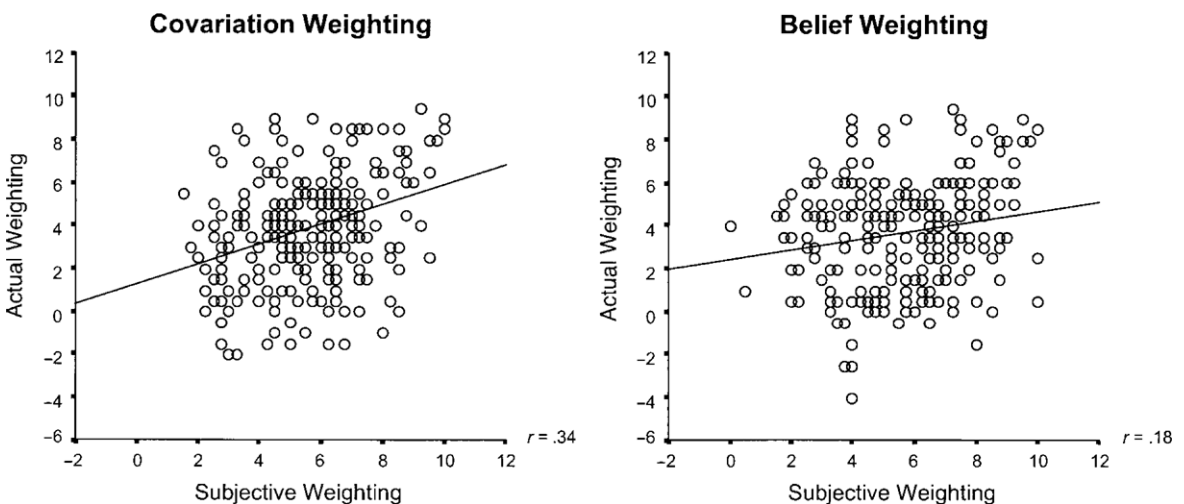


Figure 6. Experiment 3: individual calibration of participants' actual use of belief- and covariation-based cues with their subjective use of belief- and covariation-based cues for Experiment 3.

lations. Second, we observed that the reasoners were relatively accurate at articulating the extent to which their judgments were sensitive to empirical evidence but were less able to indicate the extent to which their judgments reflected their prior beliefs; again, this was the case for both familiar and unfamiliar causal relations.

Multiple Long-Term Memory Representations of Beliefs and Subsequent Biasing Effects

Three experiments were reported that demonstrated that the form of the interaction between belief level and ΔP_c depended on the nature of causal beliefs, and not just on the strength of the causal beliefs held. We have argued elsewhere (Fugelsang & Thompson, 2000) that this interaction can be thought to reflect an adaptive process wherein the search for a plausible mechanism restricts the sets of candidates for which covariation is considered. This can be thought of as a useful heuristic, given the potentially infinite number of potential causal candidates for every given effect that occurs in the natural environment. Indeed, an exhaustive analysis of all potentially covarying candidates could take an eternity. The finding that this interaction emerges only when beliefs are represented in terms of mechanism-based information suggests that the initial evaluation of naturally occurring variables represents a search for mechanistic-based evidence. If one can identify a plausible causal mechanism, covariation-based information is given more weight. The finding that this interaction is unique to causal judgments adds to the body of research that indicates that the perception of causality represents a unique cognitive capacity that is functionally dissociable from covariation estimation (e.g., Mandel & Lehman, 1998; White, 2001).

These data support the interpretation that causal beliefs are based on multiple nonredundant sources of information and argue against an interpretation in which causal beliefs are represented in a unidimensional manner. For example, according to a strict interpretation of the covariation-based view, all causal beliefs are based on the degree to which the candidate cause and the effect are observed to covary (e.g., Cheng, 1997; Rescorla, 1968; Young, 1995); this includes beliefs based on both past (i.e., long-term memory representations) and current (i.e., new empirical evidence) experience. In contrast to this view, our data suggest that people's causal beliefs reflect at least two nonredundant properties: (1) knowledge of covariation-based information and (2) knowledge of mechanism-based information.

At a minimum, therefore, models of human causal reasoning need to be modified to incorporate a role for noncovariation-based information, such as causal mechanisms, into causal beliefs. As well, these models need to explain how these causal beliefs influence the evaluation of new covariation-based evidence. One means by which covariation-based views can address the latter concern is to assume that reasoners base their judgments on a limited range of possibilities defined by a *focal set* (Cheng, 1997; Cheng & Novick, 1990, 1992).

Focal Sets and the Computation of Covariation

Cheng (1997) and Cheng and Novick (1990) argued that participants compute covariation with respect to a target population of events, which they termed a *focal set*. This target population may extend beyond the events defined by the experimenter in the problem scenario and may include additional information that is extracted from long-term memory. It is crucial, then, to know which events are deemed relevant and which irrelevant for inclusion in the focal set. Causal beliefs, regardless of how they are derived, may be a variable that sets boundaries for focal set selection.

For example, when a potential candidate is unbelievable, perhaps the relevant focal set is expanded to include instances in which the effect failed to occur when the cause was present or in which another cause produced the effect in question. When ΔP_c is computed within this larger universe of events, the presence of these alternative causes will lower the computed covariation between cause and effect. Thus, even though the nominal value of ΔP_c is high, the value that is computed within the focal set may be substantially lower than that suggested in the problem. Moreover, because this *dilution* is presumably more likely when the candidate is unbelievable than when it is believable, the effect of the presented value of ΔP_c should be smaller for unbelievable than for believable candidates.

One possible contribution of our findings might be to delimit circumstances under which reasoners either expand their focal sets beyond the range of the stimuli provided or restrict themselves to the events discussed in the problem. Therefore, our data may provide a means to rectify one of the crucial limitations of the focal set hypothesis—namely, difficulty in a priori specification (see Glymour, 2001).

Unconscious Belief-Based and Conscious Evidence-Based Processing

In Experiments 1 and 3, we observed that the reasoners were more accurate in judging the contribution of covariation-based data to their causal attributions than they were in judging the contribution of their causal beliefs. These findings are consistent with White's (1989) argument that causal beliefs may be recruited automatically and that their application to a problem-solving situation may be beyond the reasoner's conscious control. However, White (1989) provided several boundary conditions under which this automatic belief activation should occur that are not consistent with our findings. Specifically, he argued that this automatic activation of beliefs should occur only when the beliefs in question represent mechanism-based information. As such, one would not necessarily expect beliefs derived from other sources (e.g., covariation-based information) to be applied automatically. In contrast, the results of Experiment 1 suggested that the reasoners were equally poor at introspecting about the contributions of covariation-based and mechanism-based beliefs. Thus, it appears that past knowledge, regardless of how it is represented, may be applied automatically.

Our findings are also broadly consistent with the framework provided by Evans and Over (1996, 1997). They argued that rational behavior is defined in terms of an organism's ability to learn from experience and respond appropriately to a variety of familiar and novel stimuli in the environment. They proposed that most of this learning, and indeed the application of this prior knowledge, occurs unconsciously and intuitively (Evans & Over, 1996, 1997). In this respect, Evans and his colleagues (e.g., Evans, 1984, 1989, 1996; Evans & Over, 1996, 1997) have argued that much of the behavior observed in typical reasoning tasks can be attributed to unconscious processes that they describe as being both heuristic and tacit. In this model, heuristic processes determine which aspects of the problem are relevant, and the analytic processes then focus only on those aspects of the problem. Conscious processes, therefore, are thought to be engaged after the fact, so that analytic logical reasoning takes place in a circumscribed domain defined by the unconscious processes.

The findings of the present experiments are consistent with the main assumptions of this theory. Specifically, we found evidence suggesting that the recruitment of causal beliefs, a heuristic process, may occur unconsciously (or at least, with limited awareness) and, moreover, that these beliefs constrain the deliberate analytic processes needed to evaluate the empirical evidence that was provided. That is, the reasoners' evaluation of covariation-based data depended on the strength and type of belief-based information that was available. Thus, similar to Evans and Over (1996, 1997), we envision a model of reasoning that encompasses both automatic heuristic processes and deliberate analytic processes.

Finally, in a larger context, we suggest that this interaction between belief-based and empirically based processes may be an adaptive reasoning strategy. That is, when faced with an infinite number of potential causal candidates for every given effect encountered in the natural environment, it makes sense to restrict the set of candidates about which covariation information is assessed. Otherwise, an organism will expend much wasted effort examining the empirical evidence associated with implausible candidates. The practical strategy, therefore, is to import conceptual knowledge concerning the causal candidate in question in order to determine its plausibility before the decision is made to make use of any empirical cues.

A Two-Stage Model of Reasoning

In this final section, we outline a model to account for our findings and to integrate our assumptions about conscious and unconscious processes with our findings about covariation- and mechanism-based beliefs. The basic assumptions of the model are represented in Figure 7. According to our model, causal reasoning consists of two primary decision-making processes. In the first stage, individuals are assumed to recruit knowledge about the causal candidates in question. This knowledge may incorporate information about long-term memory representations of covariation-based information, long-term memory representations of mechanism-based information, or some combination of the two. This recruitment is assumed to occur automatically and with little conscious awareness on behalf of the reasoner. Stage 2 processing entails the conscious evaluation of empirical evidence (i.e., covariation-based data). Here,

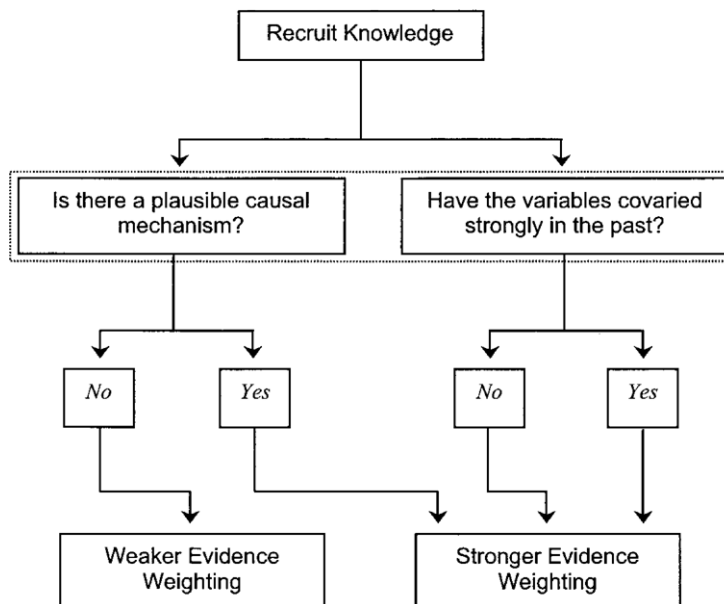


Figure 7. Graphical depiction of the proposed two-stage model of causal reasoning.

we assume that the reasoner deliberately evaluates the empirical evidence and judges the perceived causal efficacy of the candidate cause. Furthermore, along with Evans and Over (1996, 1997), we assume that the outcome of Stage 1 processing determines the course of conscious Stage 2 processing.

An explicit assumption of this model is that people represent two types of beliefs about causal relations and that these can be held independently of each other. That is, people may believe that two events covary with each other, or they may believe that a candidate cause has the capacity to produce a given effect due to the inherent properties of the events in question. These two sources of beliefs may be held in an independent or interactive fashion. That is, beliefs about long-term memory representations of covariation-based information and beliefs about causal mechanisms are nonredundant aspects of causal beliefs.

Moreover, it is assumed that beliefs about mechanism-based information have different effects on Stage 2 processing than do beliefs about covariation-based information. When covariation-based data are evaluated in light of covariation-based beliefs, the two sources of information contribute additively to the final causal judgment. In contrast, for reasons outlined above, when covariation-based data are evaluated for candidates in which a causal mechanism is available, the two sources of evidence interact, so that the covariation-based data carry more weight for believable than for unbelievable candidates.

In the present series of experiments, we have shown that the knowledge that individuals bring to bear on a task may greatly influence their ability to carry out that task in a way deemed normatively appropriate. More important, however, individuals may not be entirely aware of the extent to which their knowledge influences their decisions. These findings have implications for causal reasoning in particular and for decision making in general. For these reasons, future researchers should attempt to devise more extensive methods of accurately assessing such biases, as well as developing techniques to minimize their influence when such biases would hinder accurate decision making. In addition, because causal beliefs are certainly constructed within the cultural constraints of the social environment, a fruitful avenue for future research would be to examine cross-cultural differences in the acquisition and application of personal causal theories.

REFERENCES

- Ahn, W.-K., Kalish, C. W., Medin, D. L., & Gelman, S. A. (1995). The role of covariation versus mechanism information in causal attribution. *Cognition*, *54*, 299-352.
- Allan, L. G., & Jenkins, H. M. (1980). The judgments of contingency and the nature of response alternatives. *Canadian Journal of Psychology*, *34*, 1-11.
- Anderson, J. R., & Sheu, C.-F. (1995). Causal inferences as perceptual judgments. *Memory & Cognition*, *23*, 510-524.
- Beasley, N. E. (1968). The extent of individual differences in the perception of causality. *Canadian Journal of Psychology*, *22*, 399-407.
- Boyle, D. G. (1960). A contribution to the study of phenomenal causality. *Quarterly Journal of Experimental Psychology*, *12*, 171-179.
- Bullock, M., Gelman, R., & Baillargeon, R. (1982). The development of causal reasoning. In W. Friedman (Ed.), *The developmental psychology of time* (pp. 209-254). New York: Academic Press.
- Cheng, P. W. (1997). From covariation to causation: A causal power theory. *Psychological Review*, *104*, 367-405.
- Cheng, P. W., & Lien, Y. (1995). The role of coherence in distinguishing between genuine and spurious causes. In D. Sperber, D. Premack, & A. J. Premack (Eds.), *Causal cognition: A multidisciplinary debate* (pp. 463-494). Oxford: Oxford University Press.
- Cheng, P. W., & Novick, L. R. (1990). A probabilistic contrast model of causal induction. *Journal of Personality & Social Psychology*, *58*, 545-567.
- Cheng, P. W., & Novick, L. R. (1992). Covariation in natural causal induction. *Psychological Review*, *99*, 365-382.
- Collins, D. J., & Shanks, D. R. (2002). Momentary and integrative response strategies in causal judgment. *Memory & Cognition*, *30*, 1138-1147.
- Dennis, M. J., & Ahn, W.-K. (2001). Primacy in causal strength judgments: The effect of initial evidence for generative versus inhibitory relationships. *Memory & Cognition*, *29*, 152-164.
- Downing, C. J., Sternberg, R. J., & Ross, B. H. (1985). Multicausal inference: Evaluation of evidence in causally complex situations. *Journal of Experimental Psychology: General*, *114*, 239-263.
- Einhorn, H. J., & Hogarth, R. M. (1986). Judging probable cause. *Psychological Bulletin*, *99*, 3-19.
- Evans, J. St. B. T. (1984). Heuristic and analytic processes in reasoning. *British Journal of Psychology*, *75*, 451-468.
- Evans, J. St. B. T. (1989). *Bias in human reasoning: Causes and consequences*. Hove, U.K.: Erlbaum.
- Evans, J. St. B. T. (1996). Deciding before you think: Relevance and reasoning in the selection task. *British Journal of Psychology*, *87*, 223-240.
- Evans, J. St. B. T., Handley, S. J., & Harper, C. N. (2001). Necessity, possibility and belief: A study of syllogistic reasoning. *Quarterly Journal of Experimental Psychology*, *54A*, 935-958.
- Evans, J. St. B. T., & Over, D. E. (1996). *Rationality and reasoning*. Hove, U.K.: Psychology Press.
- Evans, J. St. B. T., & Over, D. E. (1997). Rationality in reasoning: The problem of deductive competence. *Cahiers de Psychologie Cognitive*, *16*, 3-18.
- Fugelsang, J. A., & Thompson, V. A. (2000). Strategy selection in causal reasoning: When beliefs and covariation collide. *Canadian Journal of Experimental Psychology*, *54A*, 15-32.
- Fugelsang, J. A., & Thompson, V. A. (2001). Belief-based and covariation-based cues affect causal discounting. *Canadian Journal of Experimental Psychology*, *55*, 70-76.
- Garcia, J., McGowan, B. K., Ervin, F. R., & Koelling, R. A. (1968). Cues: Their relative effectiveness as a function of the reinforcer. *Science*, *160*, 794-795.
- Glymour, C. (2001). *The mind's arrows: Bayes nets and graphical causal models in psychology*. Cambridge, MA: MIT Press.
- Harre, R., & Madden, E. H. (1975). *Causal powers: A theory of natural necessity*. Oxford: Basil Blackwell.
- Hume, D. (1978). *A treatise of human nature*. Oxford: Oxford University Press. (Original work published 1739)
- Jack, A. I., & Roepstorff, A. (2002). Introspection and cognitive brain mapping: From stimulus-response to script-report. *Trends in Cognitive Sciences*, *6*, 333-339.
- Jenkins, H., & Ward, W. (1965). Judgments of contingency between response and outcomes. *Psychological Monographs*, *7*, 1-17.
- Kant, I. (1965). *Critique of pure reason*. London: Macmillan. (Original work published 1781)
- Kao, S. F., & Wasserman, E. A. (1993). Assessment of an information integration account of contingency judgment with examination of subjective cell importance and method of presentation. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, *19*, 1363-1386.
- Klauer, K., Musch, J., & Naumer, B. (2000). On belief bias in syllogistic reasoning. *Psychological Review*, *107*, 852-884.
- Lien, Y., & Cheng, P. (2000). Distinguishing genuine from spurious causes: A coherence hypothesis. *Cognitive Psychology*, *40*, 87-137.
- Lober, K., & Shanks, D. (2000). Is causal induction based on causal

- power? Critique of Cheng (1997). *Psychological Review*, **107**, 195-212.
- Lopez, F., Shanks, D., Almaraz, J., & Fernandez, P. (1998). Effects of trial order on contingency judgments: A comparison of associative and probabilistic contrast accounts. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **24**, 672-694.
- Madden, E. H., & Humber, J. (1974). Non-logical necessity and C. J. Ducasse. In T. L. Beauchamp (Ed.), *Philosophical problems of causation* (pp. 163-178). Encino, CA: Dickenson Press.
- Mandel, D. R., & Lehman, D. R. (1998). Integration of contingency information in judgments of cause, covariation, and probability. *Journal of Experimental Psychology: General*, **127**, 269-285.
- Michotte, A. (1963). *The perception of causality*. New York: Basic Books.
- Newstead, S. E., Pollard, P., Evans, J. S. B. T., & Allen, J. (1992). The source of belief bias in syllogistic reasoning. *Cognition*, **45**, 257-284.
- Oakhill, J., & Garnham, A. (1993). On theories of belief bias in syllogistic reasoning. *Cognition*, **46**, 87-92.
- Peterson, C. R., & Beach, L. R. (1967). Man as an intuitive statistician. *Psychological Bulletin*, **68**, 29-46.
- Rescorla, R. A. (1968). Probability of shock in the presence and absence of CS in fear conditioning. *Journal of Comparative & Physiological Psychology*, **66**, 1-5.
- Salmon, W. C. (1965). The status of prior probabilities in statistical explanation. *Philosophy of Science*, **32**, 137-146.
- Schustack, M. W., & Sternberg, R. J. (1981). Evaluation of evidence in causal inference. *Journal of Experimental Psychology: General*, **110**, 101-120.
- Shultz, T. R. (1982). Rules of causal attribution. *Monographs of the Society for Research in Child Development*, **47**, 1-51.
- Shultz, T. R., Fisher, G. W., Pratt, C. C., & Rulf, S. (1986). Selection of causal rules. *Child Development*, **57**, 143-152.
- Shultz, T. R., & Ravinsky, R. B. (1977). Similarity as a principle of causal inference. *Child Development*, **28**, 1552-1558.
- Siegler, R. S., & Liebert, R. M. (1974). Effects of contiguity, regularity, and age on children's inferences. *Developmental Psychology*, **10**, 574-579.
- Spellman, B. A. (1996). Acting as intuitive scientists: Contingency judgments are made while controlling for alternative causes. *Psychological Science*, **7**, 337-342.
- Thompson, V. A., Striemer, C. L., Reikoff, R., Gunter, R. W., & Campbell, J. I. D. (2003). Syllogistic reasoning time: Disconfirmation disconfirmed. *Psychonomic Bulletin & Review*, **10**, 184-189.
- Tversky, A. (1977). Features of similarity. *Psychological Review*, **84**, 327-352.
- Tversky, A., & Kahneman, D. (1980). Causal schemas in judgments under uncertainty. In M. Fishbein (Ed.), *Progress in social psychology* (Vol. 1, pp. 49-72). Hillsdale, NJ: Erlbaum.
- Walldmann, M. R. (1996). Knowledge-based causal induction. In D. Shanks, K. Holyoak, & D. L. Medin (Eds.), *The psychology of learning and motivation: Vol. 34. Causal learning* (pp. 47-88). San Diego: Academic Press.
- Walldmann, M. R., & Hagmayer, Y. (2001). Estimating causal strength: The role of structural knowledge and processing effort. *Cognition*, **82**, 27-58.
- Wasserman, E. A., Chatlosh, D. L., & Neunaber, D. J. (1983). Perception of causal relations in humans: Factors affecting judgments of response-outcome contingencies under free-operant procedures. *Learning & Motivation*, **14**, 406-432.
- Wasserman, E. A., Dorner, W. W., & Kao, S. F. (1990). Contributions of specific cell information to judgments of interevent contingency. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **16**, 509-521.
- White, P. A. (1989). A theory of causal processing. *British Journal of Psychology*, **80**, 431-454.
- White, P. A. (1992). Causal powers, causal questions, and the place of covariation information in causal attribution. *British Journal of Psychology*, **83**, 161-188.
- White, P. A. (1995). Use of prior beliefs in the assignment of causal roles: Causal powers versus regularity-based accounts. *Memory & Cognition*, **23**, 243-254.
- White, P. A. (2000). Causal judgment from contingency information: The interpretation of factors common to all instances. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **26**, 1083-1102.
- White, P. A. (2001). Causal judgments about relations between multi-level variables. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **27**, 499-513.
- White, P. A. (2002a). Causal judgment from contingency information: Judging interactions between two causal candidates. *Quarterly Journal of Experimental Psychology*, **55A**, 819-838.
- White, P. A. (2002b). Perceiving a strong causal relation in a weak contingency: Further investigation of the evidential evaluation model of causal judgment. *Quarterly Journal of Experimental Psychology*, **55A**, 97-114.
- Young, M. E. (1995). On the origin of personal causal theories. *Psychonomic Bulletin & Review*, **2**, 83-104.

APPENDIX

Causal Scenarios Used in Experiments 1 and 2

The six causal scenarios are presented below. Note that all six scenarios were used in Experiment 2 and only four of the scenarios (*Cancer*, *Fatigue*, *Flowers Blooming*, and *Slippery Roads*) were used in Experiment 1. All the scenarios were presented equally often, using both of the covariation contingencies; however, only the highly contingent ($\Delta P_c = .9$) versions are illustrated here. In addition, for the sake of brevity, the scenarios are presented only in the format utilized in Experiment 2.

Fatigue

Imagine you are a doctor trying to determine the cause of fatigue in a group of patients. You have a hypothesis that the fatigue may be due to insomnia. In order to test this theory, you investigate 10 patients that had insomnia and 10 patients that did not have insomnia. A thorough investigation revealed the following information: of the 10 patients that had insomnia, 9 were fatigued; of the 10 patients that did not have insomnia 0 were fatigued.

Slippery Roads

Imagine you are a city researcher who is trying to determine the cause of slippery roads in a group of townships. You have a hypothesis that the slippery roads may be due to ice storms. In order to test this theory, you investigate 10 townships that had ice storms and 10 townships that did not have ice storms. A thorough investigation revealed the following information: of the 10 townships that had ice storms, 9 experienced slippery roads; of the 10 townships that did not have ice storms, 0 experienced slippery roads.

Flowers Blooming

Imagine you are a horticulturist who is trying to determine the cause of flowers blooming. You have a hypothesis that the flowers blooming may be due to being planted in red pots. In order to test this theory, you investigate 10 flowers that were planted in red pots and 10 flowers that were not planted in red pots. A thorough investigation revealed the following informa-

APPENDIX (Continued)

tion: of the 10 flowers that were planted in red pots, 9 bloomed; of the 10 flowers that were not planted in red pots, 0 bloomed.

Cancer

Imagine you are a scientist who is trying to determine the cause of the development of lung cancer in a group of patients. You have a hypothesis that the lung cancer may be due to taking iron supplements. In order to test this theory, you investigate 10 patients who took iron supplements and 10 patients who did not take iron supplements. A thorough investigation revealed the following information: of the 10 patients who took iron supplements, 9 had lung cancer; of the 10 patients who did not take iron supplements, 0 had lung cancer.

Fever

Imagine you are a school nurse who is trying to determine the cause of a recent surge in fevers among children. You have

a hypothesis that the fevers may be due to having chills. In order to test this theory, you investigate 10 children who had chills and 10 children who did not have chills. A thorough investigation revealed the following information: of the 10 children who had chills, 9 had fevers; of the 10 children who did not have chills, 0 had fevers.

Exam Success

Imagine you are a professor who is trying to determine the cause of exam success. You have a hypothesis that the exam success may be due to paper writing success. In order to test this theory, you investigate 10 students who had paper writing success and 10 students who did not have paper writing success. A thorough investigation revealed the following information: of the 10 students who had paper writing success, 9 wrote successful exams; of the 10 students who did not have paper writing success, 1 wrote a successful exam.

(Manuscript received March 17, 2003;
accepted for publication March 24, 2003.)