

Conditional reasoning and causation

DENISE D. CUMMINS, TODD LUBART, OLAF ALKSNIS, and ROBERT RIST
Yale University, New Haven, Connecticut

An experiment was conducted to investigate the relative contributions of syntactic form and content to conditional reasoning. The content domain chosen was that of causation. Conditional statements that described causal relationships (if <cause>, then <effect>) were embedded in simple arguments whose entailments are governed by the rules of truth-functional logic (i.e., *modus ponens*, *modus tollens*, *denying the antecedent*, and *affirming the consequent*). The causal statements differed in terms of the number of alternative causes and disabling conditions that characterized the causal relationship. (A disabling condition is an event that prevents an effect from occurring even though a relevant cause is present.) Subjects were required to judge whether or not each argument's conclusion could be accepted. Judgments were found to vary systematically with the number of alternative causes and disabling conditions. Conclusions of arguments based on conditionals with few alternative causes or disabling conditions were found to be more acceptable than conclusions based on those with many.

A characteristic of human reasoning performance is that it is influenced by the content, or subject matter, of the reasoning task. Reasoning problems with identical formal properties but different subjective contents often produce different levels of performance. Such *content effects* have been reported in a variety of domains, including causal attribution tasks (Cheng & Novick, 1990; Novick, 1990), categorical syllogisms (Evans, 1982; Evans, Barston, & Pollard, 1983; Revlin, Leirer, Yopp, & Yopp, 1980; Wilkins, 1928), and conditional reasoning tasks (Cheng & Holyoak, 1985; Cheng, Holyoak, Nisbett, & Oliver, 1986; Griggs & Cox, 1982; Manktelow & Evans, 1979; Reich & Ruth, 1982; Wason & Johnson-Laird, 1972). Conditional reasoning provides one of the richest sources of content effects in human reasoning due to the multitude of conversational implicatures conditional statements are used to convey (e.g., threats, promises, predictions, permissions, causation). Consider, for example, the following two arguments:

1. If it is a dog, then it is an animal.
It is not a dog.
Therefore, it is not an animal.
2. If you mow the lawn, then I will give you \$5.
You do not mow the lawn.
Therefore, I will not give you \$5.

Although both arguments are of the same (invalid) form (*denying the antecedent*), people generally find Argument 2 but not 1 acceptable (Fillenbaum, 1975, 1976, 1978). This is because our prior knowledge concerning

promises indicates that conditional statements like that in Argument 2 are to be interpreted as biconditionals; that is, if you mow the lawn, I will give you \$5, *and* if you do *not* mow the lawn, then I will *not* give you \$5. Hence, Argument 2 appears valid.

An important question concerning these content effects is what they imply about the quality and soundness of "natural" human reasoning strategies. As Henle (1962) observed, when reasoning deductively, people often "refuse to accept the logical task"; that is, they refuse to reason from the stated premises alone, choosing instead to add, delete, or modify premises on the basis of their knowledge of the content domain, as in the above examples. Although this propensity to modify premises on the basis of prior knowledge is often considered a cause of reasoning errors in the deductive reasoning literature (see Evans, 1989, for a review), knowledge-based modification of problem statements lies at the heart of expert problem solving, where statements about (e.g.) inclined planes are augmented with appropriate knowledge concerning theoretical entities and principles in order to derive a solution (Bassok, 1990; Bassok & Holyoak, 1989; Chi, Feltovich, & Glaser, 1981; Larkin, 1981; Larkin, McDermott, Simon, & Simon, 1980). In fact, it is the nature and quality of this imported knowledge that often determines solution success.

Theories of human reasoning, however, differ in their treatment of such pragmatically driven content effects. Formal theorists make a clear distinction between *deductive inference* and *pragmatic inference*—that is, between reasoning on the basis of form alone (deductive inference) and reasoning on the basis of content alone or some mixture of the two (pragmatic inference) (Braine, 1978; Braine, Reiser, & Rumin, 1984; Rips, 1983; Rumin, Connell, & Braine, 1983). The belief here is that beneath pragmatic influences lies a deductive inference system that, like formal logical systems, is driven solely by the

We would like to thank Patricia Cheng, Jonathon St. B. Evans, Philip Johnson-Laird, Alvin Goldman, Robert Cummins, and David Sanford for helpful comments on an earlier version of this manuscript. Requests for reprints should be sent to Denise Cummins, Psychology Department, University of Arizona, Tucson, AZ 85721.

syntactic form of the premises and arguments. The foundations and rules upon which the human reasoning system rests may differ from those of formal logic, but in both systems, inferences are governed by form. For example, in their theory of natural logic, Braine and his colleagues propose that the natural interpretation assigned to conditional statements, regardless of content, is simply that one can conclude q given that p is known to be true (see Braine, 1978, p. 8). Reasoning performance that indicates a deviation from this interpretation is attributed to an *expansion* of the premise set to include conversational implicatures or other additional premises that represent relevant factual information from the reasoner's knowledge base. Such expansion is often considered an error in the interpretation stage of the reasoning process:

The theory allows three sources for reasoning errors. . . . A comprehension error is an error of construal of the premises or of the conclusion: The starting information used by the subject is not that intended by the problem setter. . . . These comprehension errors frequently stem from the fact that "premises might have 'conversational implicatures'" (Grice, 1975) that would lead to error. Thus it is well known that a conditional (*if p then q*) invites the inference *If not p then not q* (Geis & Zwicky, 1971), which leads to the standard fallacies of conditional reasoning. (Braine et al., 1984, p. 318)

Braine and his colleagues go on to state that predicting pragmatic inference would require "parameters estimating the effective premise set" (Braine et al., 1984, p. 360). In other words, in order to reliably predict pragmatic influences on reasoning, what is needed is some way of predicting a priori (1) the types of premises that are likely to be imported into the reasoning task as a function of content, and (2) how these premises will influence reasoning.

Such an account was recently offered by Cheng and her colleagues; in it, content determines the semantics and entailments of conditionals via pragmatic reasoning schemas (Cheng & Holyoak, 1985; Cheng et al., 1986). A pragmatic reasoning schema is a body of rules induced from life experiences that captures the proper interpretation of statements referring to classes of situations, such as permission situations, obligations, threats, and so forth. These interpretations do not necessarily correspond to those assigned to the logical (material) conditional in first-order logic. Nor should they. Cheng et al. argue this point most strongly by showing that people often import modals ("must," "can," etc.) when rephrasing certain conditionals (e.g., permissions), and modals are beyond the scope of truth-functional logic. From this view, therefore, human reasoning is characterized primarily as pragmatic inference, and variations in interpretations due to content are outcomes of this inferential process.

Recent work indicates that one major pragmatic influence on conditional reasoning is awareness of possible alternatives to the antecedent condition. For example, adding the premise "If you wash the dishes, I will give you \$5" to Argument 2 above suppresses the biconditional interpretation of the first conditional, since it sug-

gests other conditions that may lead to my giving you \$5 other than your mowing the lawn (Byrne, 1989; Romain et al., 1983). More importantly, such alternatives need not be explicitly presented to alter interpretations of conditional premises. Markovits (1986) presented subjects with logical arguments based on conditionals with familiar causal content (e.g., "If one eats poisoned mushrooms, then one becomes ill"), and less familiar causal content (e.g., "If a proton strikes a neutron, then the temperature rises"). He also required subjects to generate at least one counterexample for the two conditionals (e.g., cases in which one becomes ill, but one has *not* eaten poisoned mushrooms). Although subjects were able to generate at least one alternative for each conditional, reasoning performance on the familiar content conditional exceeded that on the less familiar one, indicating that a conditional interpretation was more likely for the familiar than for the less familiar case. On the basis of this, Markovits concluded that familiarity influences reasoning by increasing awareness that such counterexamples exist, regardless of whether such counterexamples are actually generated.

Importantly, Markovits (1986) did not report the number of counterexamples, or alternatives, generated in the familiar and less familiar cases. It seems reasonable to assume that subjects could generate more alternatives for the familiar situation than for the unfamiliar one. If this were the case, an alternative explanation for his results could be that conditionals for which a number of counterexamples are possible are interpreted differently from those for which only a few are considered possible. For example, the tendency to interpret a statement as a conditional or a biconditional may exist on a continuum, varying with the size of the pool of alternatives that characterize the situation described by the conditional.

Another difficulty in interpreting Markovits's (1986) results is the possibility of an interaction between the nature of the alternative and the form of the argument. Byrne (1989), for example, found that certain arguments were sensitive to the type of alternative information stated in added premises. In the case of causal conditionals, there are two types of possible alternatives: possible causes that could produce the effect in question, and possible disabling conditions that could prevent the effect from occurring despite the presence of a cause. To see why the nature of the alternative is important, consider the following:

If my finger is cut, then it bleeds.

My finger is cut.

Therefore, it bleeds.

If I eat candy often, then I have cavities.

I eat candy often.

Therefore, I have cavities.

Both arguments are of exactly the same form, namely *modus ponens*; hence both are deductively valid. Most people, however, would probably be more inclined to accept the conclusion of the first argument than that of the second because one can think of few things that would prevent a finger from bleeding if it were cut, but many

things that would prevent someone from getting cavities even if candy were eaten. This suggests that performance on *modus ponens* should be sensitive to disabling conditions. Now consider the following:

If my finger is cut, then it bleeds.
My finger is bleeding.
Therefore, my finger is cut.

If I eat candy often, then I have cavities.
I have cavities.
Therefore, I eat candy often.

Both of these arguments are of the form *affirming the consequent*; hence both are deductively invalid. Yet, again, most people would probably be more inclined to accept the first argument than the second. This is because one can think of few things that could cause a finger to bleed other than cutting the skin on it, but one can think of many things that could cause someone to have cavities other than eating candy. Thus, while *modus ponens* should be influenced by possible disabling conditions, *affirming the consequent* should be influenced by possible alternative causes. The same can be argued for two other types of commonly studied arguments—namely, *modus tollens* and *denying the antecedent*. Like *modus ponens*, *modus tollens* should be influenced by alternative disabling conditions:

If my finger is cut, then it bleeds.
My finger is *not* bleeding.
Therefore, my finger is *not* cut.

If I eat candy often, then I have cavities.
I do *not* have cavities.
Therefore, I do *not* eat candy often.

In contrast, *denying the antecedent*, like *affirming the consequent*, should be influenced by alternative causes:

If my finger is cut, then it bleeds.
My finger is *not* cut.
Therefore, it is *not* bleeding.

If I eat candy often, then I have cavities.
I do *not* eat candy often.
Therefore, I do *not* have cavities.

To summarize, we predicted that prior knowledge concerning alternative possibilities to the causal scenarios described by causal conditionals would influence peoples' interpretations of the conditionals, and hence their willingness to accept or reject conclusions based on them. More particularly, we predicted an interaction between the nature of the prior knowledge and the form of the argument on reasoning performance. *Modus ponens* and *modus tollens* were predicted to be influenced by the number of disabling conditions that characterized their embedded causal conditionals, while *affirming the consequent* and *denying the antecedent* were predicted to be influenced by the number of alternative causes characterizing the conditionals. These hypotheses are depicted in Figure 1. Empirical support for these hypotheses would provide evidence concerning (1) the type of pragmatic knowledge that

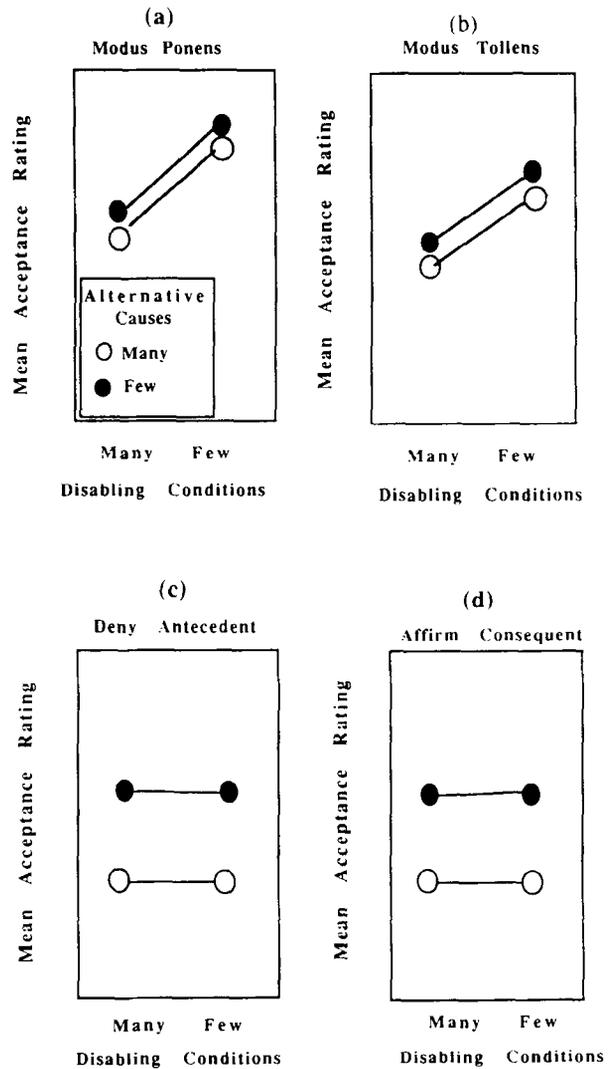


Figure 1. Predicted performance differences on arguments containing causal conditionals that have many or few possible causes and many or few possible disabling conditions. Relative performance on the first two argument types (a and b) is predicted to be influenced by the number of possible disabling conditions possessed by their embedded conditionals; performance on the latter two (c and d) is predicted to be influenced by the number of possible causes.

is considered during reasoning with causal conditional statements, and (2) precisely how this knowledge influences interpretation of the conditional statement and, hence, the conditional reasoning process.

METHOD

Subjects

Twenty-seven students who were enrolled in introductory psychology during the summer session at Yale University served as subjects in the experiment. None had studied logic before.

Materials

A summary of the experimental materials and design is presented in Figure 2. Conditional statements for the tasks were chosen in

the following way: A pool of 54 causal conditional statements was created, by gathering them from a variety of sources. This pool of statements was presented to a different group of 22 subjects (introductory psychology students during the previous spring term at Yale). One half of these subjects were required to generate possible causes for the effects stated in the conditional, and one half were required to generate possible disabling conditions. An example of the cause-generation task is the following:

Rule: If Joyce eats candy often, then she will have cavities.
 Fact: Joyce has cavities, but she does *not* eat candy often.

Please write down as many circumstances as you can that could make this situation possible.

An example of the disabling conditions generation task is:

Rule: If Joyce eats candy often, then she will have cavities.
 Fact: Joyce eats candy often, but she does *not* have cavities.

Please write down as many circumstances as you can that could make this situation possible.

Formats like these were constructed for each of the 54 conditionals; they were typed one to a page in a booklet. The subjects were allowed 1.5 min to respond to each conditional statement during the generation task. (The time allowed was chosen on the basis of pilot work. Subjects generally tended to "run dry" after approximately 1 min.) Following this, all subjects were given another booklet containing the 54 conditionals, and they were asked to rate the strength of the causal relationship described by each. They did this by choosing a number from a 6-point rating scale in which 0 stood for "no causal relationship" and 5 stood for "very strong causal relationship." It was explained that a cause is an event that produces another event, and that the subjects' ratings should reflect the likelihood that the event in the "if" part of the statement could produce the event in the "then" part of the statement. This task was used to ensure that the statements we chose were in fact describing causal relationships from the subjects' point of view.

The 54 conditionals were rank-ordered in terms of the mean rating given on the causal strength task from strong to weak. The top 28 conditionals (i.e., those that subjects believed to be describing

strong causal relationships) were chosen for further analyses. The mean rating for the 28th conditional was 3.36, which corresponded to "moderately strong causal relationship" on the rating scale.

Generation protocols were scored by two independent raters, using scoring criteria that were decided on in advance. The main point of these criteria was to disallow unlikely instances (e.g., aliens intervened) and to disallow multiple instances that were variants of a single idea. Using the cause-generation example above, legal instances would be (e.g.) brushing after eating candy, having frequent fluoride treatments by a dentist, and genetic makeup. Illegal instances would be (e.g.) brushing after eating candy, brushing after meals, brushing several times a day. In this case, all the instances are variants on a single theme: brushing. As a result, these would count as only one cause. Very few such instances (less than 4%) were observed, presumably because the instructions discouraged the listing of simple theme variations.

From these data, the mean number of causes and disabling conditions for each of the remaining 28 conditional statements was calculated, along with standard deviations. These means were then divided into quartiles, and from these quartiles, four classes of stimuli were chosen. The first included four statements that appeared in the upper quartile for both number of causes and number of disabling conditions; the second included four that appeared in the bottom quartile on both; the third included four that appeared in the upper quartile for causes and in the bottom quartile for disabling conditions; and the fourth included four that appeared in the bottom quartile for causes and in the upper quartile for disabling conditions. In each case, the four chosen from within the quartiles were those that had the lowest standard deviations (and hence most subject agreement). The 16 chosen statements therefore represented four instances within a 2 x 2 matrix of alternative causes (many or few) and disabling conditions (many or few). The mean number of alternative causes generated and mean causal strength ratings within each cell are presented in the bottom half of Figure 2.

The chosen 16 conditionals were then embedded in the four types of arguments—that is, *modus ponens* (MP), *modus tollens* (MT), *denying the antecedent* (DA), and *affirming the consequent* (AC), thereby producing 64 arguments. These 64 were bound into four booklets, with each booklet containing one instance from each of the cause x disabling condition x argument type conditions. Present-

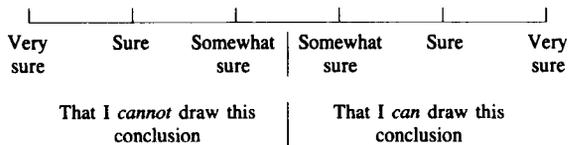
Alternative Causes	Many								Few							
	Many				Few				Many				Few			
Disabling Conditions	Many				Few				Many				Few			
Argument Type	MP	MT	DA	AC	MP	MT	DA	AC	MP	MT	DA	AC	MP	MT	DA	AC
Conditional	C1	C1	C1	C1	C5	C5	C5	C5	C9	C9	C9	C9	C13	C13	C13	C13
	C2	C2	C2	C2	C6	C6	C6	C6	C10	C10	C10	C10	C14	C14	C14	C14
	C3	C3	C3	C3	C7	C7	C7	C7	C11	C11	C11	C11	C15	C15	C15	C15
	C4	C4	C4	C4	C8	C8	C8	C8	C12	C12	C12	C12	C16	C16	C16	C16
	Mean		Std. Dev.		Mean		Std. Dev.		Mean		Std. Dev.		Mean		Std. Dev.	
Alternative Causes	3.57		2.35		3.75		1.89		2.50		1.93		2.16		1.75	
Disabling Conditions	3.66		2.27		2.77		1.79		3.84		1.99		2.04		2.14	
Causal Ratings	4.30		1.63		4.00		1.76		3.88		1.96		4.56		1.18	

Figure 2. Diagram of experimental design and description of stimulus materials used in the experiment. The top half of the diagram depicts the materials included in the booklets given to subjects. Each booklet contained 64 arguments, 4 each based on 16 conditional statements. The conditionals differed in terms of the mean number of alternative causes (many or few) and the mean number of disabling conditions (many or few) that characterized them. The four types of arguments were *modus ponens* (MP), *modus tollens* (MT), *denying the antecedent* (DA), and *affirming the consequent* (AC). The bottom half of the figure presents the mean number of alternative causes and disabling conditions and mean causal strength ratings generated by a pilot group of subjects for the stimuli in each alternative causes x disabling condition cell.

tation order of arguments within booklets was randomized, and presentation order of the booklets was counterbalanced across subjects.

Each page in the booklets contained one argument. Beneath the argument was a rating scale on which subjects were to record their judgment concerning the argument's conclusion. The scale was as follows:

Given this rule and this fact, place a mark on the scale below that best reflects your evaluation of the conclusion



To summarize: stimulus materials consisted of 64 arguments, 4 instances within each of the 16 argument type (MP, MT, DA, AC) \times causes (few, many) \times disabling conditions (few, many) cells. Each subject therefore contributed four observations in each cell. The subjects were told that their task was to *determine whether the conclusion could be drawn from the stated premises*.

Procedure

The subjects were run as a group. The booklets were distributed among them, along with written instructions, one example, and one practice argument. They were also orally instructed in the nature of the task; the oral instructions included a reminder to work the problems in the order in which they appeared in the booklets. The subjects were *not* instructed about the "proper" interpretation of conditional statements, for two reasons: First, we were interested in the natural interpretations that people apply to such statements, not in how well they could learn the rules of first-order logic. Second, although the proper interpretation of abstract (material) conditionals is specified by truth-functional logic, the proper interpretation of causal conditionals is still subject to debate. The subjects also were not explicitly told to assume that the premises were true. Instead, we allowed the contextual cues "Rule" and "Fact" to suggest this. Again, we were interested in observing human reasoning performance as it typically occurs, and not in instructing subjects in deductive logic.

RESULTS

Rejection probability for all statistical tests was .05, unless noted otherwise. Significant interactions were further analyzed with simple effects tests (Keppel, 1973); significant main effects and simple main effects involving more than two means were analyzed with Tukey's *HSD* test of pairwise comparisons.

Acceptance Ratings

Rating choices were mapped onto a 6-point scale from -3 to $+3$, with -3 corresponding to "very sure *cannot* draw this conclusion" and $+3$ corresponding to "very sure *can* draw this conclusion." Subjects' ratings for the four instances within each argument \times causes \times disabling condition cells were averaged. Cell means are depicted in Figures 3a-3d. The first thing to notice is that all means were positive, indicating that, on the average, subjects tended to find all four arguments based on these causal conditionals acceptable, suggesting a biconditional interpretation. The degree of acceptability varied greatly, however.

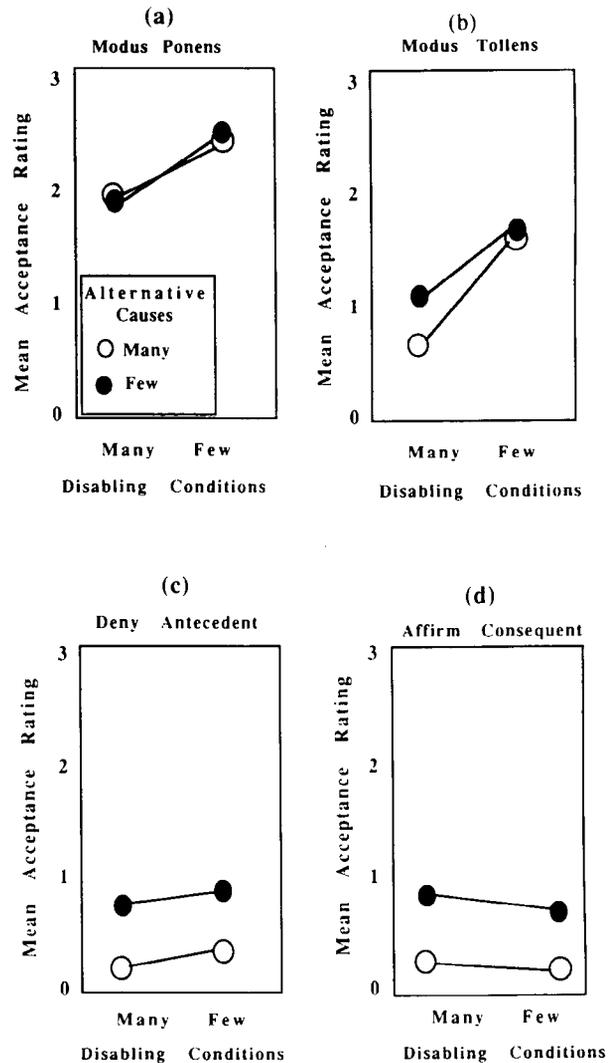


Figure 3. Observed mean acceptance ratings for each type of argument. Positive scores indicate acceptance of the conclusions and negative scores indicate rejections.

An analysis of variance (ANOVA) was performed on the mean acceptance ratings, using as variables argument (MP, MT, DA, and AC), number of causes (many and few), and number of disabling conditions (many and few). The main effects of argument, number of causes, and number of disabling conditions were all significant [$F(3,78) = 22.65$, $MS_e = 2.74$, $p < .001$; $F(1,26) = 18.07$, $MS_e = .66$, $p < .01$; and $F(1,26) = 18.03$, $MS_e = .53$, $p < .01$, respectively]. These main effects were modified by higher order interactions [$F(3,78) = 4.36$, $MS_e = .30$, $p < .01$, and $F(3,78) = 12.27$, $MS_e = .28$, $p < .01$, for the argument \times causes and arguments \times disabling conditions interactions, respectively]. As predicted, ratings for *modus ponens* and *modus tollens* were influenced by disabling conditions [$F(1,26) = 16.07$ and 42.61 , $MS_e = .21$ and $.36$, respectively]. As is apparent in Figures 3a and 3b, subjects gave lower acceptance ratings to these arguments when they were based on conditionals characterized by

many disabling conditions than when the arguments were based on conditionals with few disabling conditions. Also as predicted, disabling conditions did *not* influence ratings for *denying the antecedent* or *affirming the consequent* [$F_s(1,26) = 1.92$ and $.73$, $MS_e_s = .46$ and $.35$, $ps > .05$, respectively].

Turning now to the influence of alternative causes, our predictions were again supported, with one modification. As predicted, ratings for *denying the antecedent* and *affirming the consequent* were only influenced by number of alternative causes [$F_s(1,26) = 9.23$ and 21.33 , $MS_e_s = .69$ and $.37$, $ps < .01$, respectively]. Contrary to predictions, however, *modus tollens* was also found to be influenced by alternative causes [$F(1,26) = 5.76$, $MS_e = .21$, $p < .05$]. Nonetheless, the nature of this influence was the same for all three argument types: Ratings moved toward acceptance when the arguments described effects characterized by few alternative causes, and toward rejection when the argument described effects characterized by many alternative causes.

Influence of Alternatives on Interpretation of the Conditional

As stated earlier, the presence of *few* alternatives may force a biconditional interpretation of causal conditional statements. Indeed, even formal theorists maintain that the biconditional interpretation may be warranted for certain conditionals, notably those describing temporal-causal situations (Marcus & Rips, 1979). The types of interpretations subjects assigned to the conditionals can be inferred from their response patterns. Consider the examples given earlier for the conditional "If I cut my finger, then it bleeds." Embedding this conditional in the four types of arguments produces the following:

- MP: If my finger is cut, then it bleeds.
My finger is cut.
Therefore, it bleeds.
- MT: If my finger is cut, then it bleeds.
My finger is *not* bleeding.
Therefore, my finger is *not* cut.
- DA: If my finger is cut, then it bleeds.
My finger is *not* cut.
Therefore, it is *not* bleeding.
- AC: If my finger is cut, then it bleeds.
My finger is bleeding.
Therefore, my finger is cut.

If subjects interpreted this if-then statement as a biconditional, they would have *accepted* the conclusions of all four arguments. If acceptances are coded as As and rejections as Rs, then a biconditional response pattern for this conditional would be AAAA, meaning that the subject accepted all four arguments when they were based on this conditional. If subjects interpreted this if-then statement as a material conditional, however, they would have accepted the conclusions of MP and MT, and rejected the conclusions of DA and AC, producing a

Table 1
Mean Proportions of Response Patterns That Were Consistent with a Biconditional or Conditional Interpretation

	Disabling Conditions			
	Many		Few	
	Many Causes	Few Causes	Many Causes	Few Causes
Biconditional	.18	.41	.29	.36
Conditional	.16	.15	.26	.27

Note—Mean proportions based on four observations from each of 27 subjects.

response pattern of AARR for this conditional. Note that the major difference between the two interpretations is how DA and AC are evaluated. Since the judgment task results clearly showed that these arguments are particularly sensitive to alternative causes, we would expect to find more biconditional interpretations (i.e., AAAA) for conditionals that are characterized by few causes than for those that are characterized by many.

To test this possibility, response patterns were coded and classified into three categories: conditional interpretations (AARR), biconditional interpretations (AAAA), or "other" interpretation (i.e., any deviation from these two patterns). Positive acceptance ratings were considered acceptances, and negative ratings were considered rejections. The biconditional and conditional interpretations together accounted for 52% of all response patterns observed. (Ninety-two percent of all subjects evidenced at least one biconditional response pattern, and 70% evidenced at least one conditional response pattern.)

As is apparent from Figure 2, each subject provided four observations within each 2 (alternative causes) × 2 (disabling condition) cell. The proportion of conditional and biconditional interpretations produced by each subject within each cell was calculated. Mean proportions are reported in Table 1. An ANOVA was performed on these data, using three within-subjects variables: number of causes (many and few), number of disabling conditions (many and few), and response pattern (biconditional and conditional). (Responses in the "other" category were excluded from the analysis, since there are only two degrees of freedom when partitioning responses into three categories.) The main effects of causes and disabling conditions were significant [$F_s(1,26) = 16.22$ and 10.35 , $MS_e_s = .02$ and $.03$, $ps < .01$, respectively]. In the case of disabling conditions, the direction of the means indicated that more biconditional and conditional interpretations occurred when few disabling conditions were possible than when many were possible.

More importantly, the response pattern × causes interaction was also significant [$F(1,26) = 8.41$, $MS_e = .04$, $p < .01$]. This interaction is depicted in Figure 4. The interpretation is straightforward; as predicted, biconditional interpretations were more likely among conditionals with few alternative causes than they were among those with many [$F(1,26) = 19.91$, $MS_e = .03$, $p < .001$]. Conditional interpretations were unaffected by the number of

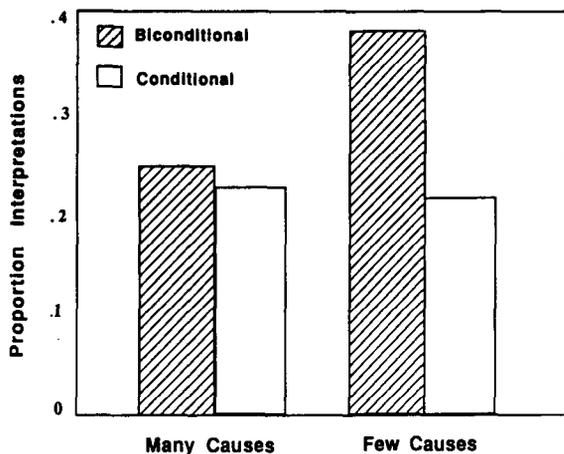


Figure 4. Mean proportion response patterns that were consistent with biconditional or conditional interpretations as a function of the number of alternative causes that characterized the causal conditional statement.

alternative causes ($F < 1$). There was also some tendency for biconditional interpretations to be preferred over conditional interpretations when few alternative causes were possible, but the difference was only marginally significant [$F(1,26) = 3.69$, $MS_e = .23$, $.05 < p < .07$].

These results suggest that prior knowledge concerning alternatives influences conditional reasoning by influencing the interpretation assigned to the conditional statement. When the situation described by a causal conditional statement suggests few disabling conditions, the statement is likely to be assigned either a conditional or a biconditional interpretation (out of the numerous possible interpretations that may be assigned). If the situation also suggests few alternative causes, the statement is most likely to be assigned a biconditional interpretation. Most importantly, these results provide clear evidence concerning the type of prior knowledge that people are likely to access when interpreting causal conditionals (i.e., alternative causes and disabling conditions) and precisely how that accessed knowledge affects their interpretations.

DISCUSSION

The pattern of results reported here is consistent with those of a large body of literature showing that conditional reasoning is influenced by the content of the conditional. More particularly, this pattern suggests that the interpretation people assign to conditional statements is influenced by the size of the pool of possible alternatives for the scenarios described by conditional statements. In the case of causal conditionals, these alternatives take the form of alternative causes and disabling conditions. Together, these results strongly suggest that the interpretations assigned to conditionals are sensitive not just to the type of predication or conversational implication they suggest (e.g., temporal/causal, permission, threat, etc.), but to

other knowledge-based factors as well. The most important of these appear to be the nature and number of alternatives or counterexamples they suggest.

Consideration of alternatives appears to be a cornerstone in many aspects of human cognition. Theories that appeal to alternatives include the mental model theory of reasoning offered by Johnson-Laird (1983), the norm theory of judgment and decision-making proposed by Kahneman and Miller (1986), the exemplar theory of classification learning offered by Medin and his colleagues (e.g., Medin, 1983; Medin, Altom, Edelson, & Freko, 1982; Medin & Schaffer, 1978), and the causal attribution theories of Kelley (1972, 1973) and Cheng and Novick (1990). According to all of these theories, a stimulus selectively recruits its own alternatives, and these alternatives provide a rich context of remembered and constructed representations with which to interpret or classify the stimulus. Performance is influenced by the number of available alternatives or category exemplars. In the case of judgment and decision-making, these alternatives can be counterfactuals that are constructed ad hoc on the basis of general knowledge rather than specific situations retrieved from past experience. The work presented here is consistent with this approach, indicating that searching for alternatives is a general reasoning strategy that people often use when performing cognitive tasks.

The work of Johnson-Laird (1983) on mental models is particularly relevant here. From this view, a conditional describes a possible state of affairs that one represents internally as a model. Reasoning consists of a strategy of searching memory for counterexamples to the models constructed. In a sense, a mental model constructed from a conditional is a situation-specific representation of the state of the world suggested by the propositions contained in the conditional, and reasoning consists of searching for alternatives to the situation. Yet although a "search for counterexamples" may be a general reasoning strategy, the *nature* of the counterexamples and the role they play in reasoning may differ in terms of (and hence be reliably predicted from) the content of the conditional. The work reported here suggests that, in the case of causal reasoning, alternatives take the form of alternative causes and disabling conditions. Counterexamples or alternatives in other domains would presumably be of different types. More importantly, the impact of such alternatives may be very different across content areas. For example, whereas discovering a nonflying bird may cause us to reject the conditional "If it is a bird, then it flies," and hence reorganize our conceptual framework for birds, discovering that a particular effect (call it "Effect B") did not follow a hypothesized Cause A may not lead us to reject the conditional "If Cause A is present, then Effect B occurs." Instead, background assumptions concerning the presence of disabling conditions may be called into question. The types of entailments that are assigned to the conditionals in the two content areas may differ systematically as well. More research is needed to identify the nature of the al-

ternatives considered in different content domains, as well as how these alternatives affect the interpretation and entailments assigned to the conditionals.

Also of particular relevance here is the *discounting principle* suggested by Kelley (1972, 1973), which states that the role of any particular cause in producing an effect is discounted if other possible causes or intervening circumstances are present. Kelley introduced the discounting principle to account for human performance patterns on causal attribution tasks—that is, tasks in which one has to determine which of several candidate causes actually produced an event. These results strongly suggest, however, that the discounting principle also applies to causal conditional reasoning performance. Here, discounting produces a disinclination on the part of the reasoner to accept the conclusions of arguments based on conditional scenarios with many alternatives.

It should also be noted, however, that content is not the whole story. Performance on *modus ponens* was consistently higher than any other argument in all conditions, a result that is consistent with Braine's (1978) formal theory of natural deduction in which *modus ponens* constitutes the "natural" interpretation assigned to conditional statements (that is, "if p, then q" is interpreted to mean, "If p is asserted as true, then one is allowed to conclude q"). Yet even this argument was influenced by prior knowledge. These results are consistent with those of Evans et al. (1983), which showed that syllogistic reasoning is influenced by both argument form and prior beliefs, and that belief "bias" is more marked on invalid than on valid forms. Clearly, people are sensitive to syntactic form when reasoning deductively. It is also important to note that Evans et al.'s subjects were explicitly instructed to make deductive inferences, but that even here, pragmatic considerations were evident in their performance.

Finally, this work underscores the point made by others in the field (e.g., Braine, 1978; Braine et al., 1984; Henle, 1962) regarding the importance of ensuring that one has captured the correct formal representation of a natural language argument as represented by the reasoner when evaluating human reasoning performance. There is abundant evidence that people's representations of the truth conditions for certain logical statements go beyond those of first-order logic, and that the knowledge they import into their representations of the reasoning problem is often relevant to the task of drawing true conclusions to guide behavior. More importantly, however, this work strongly suggests that in the case of causal conditional reasoning, interpretation of the conditional statement depends (at least in part) on the size of the set of alternative causes and disabling conditions that characterize the cause-effect relation in question.

REFERENCES

BASSOK, M. (1990). Transfer of domain-specific problem-solving procedures. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **16**, 522-533.
 BASSOK, M., & HOLYOAK, K. J. (1989). Interdomain transfer between

isomorphic topics in algebra and physics. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **15**, 153-166.
 BRAINE, M. D. S. (1978). On the relation between the natural logic of reasoning and standard logic. *Psychological Review*, **85**, 1-21.
 BRAINE, M. D. S., REISER, B. J., & RUMAIN, B. (1984). Some empirical justification for a theory of natural propositional logic. In G. H. Bower (Ed.), *The psychology of learning and motivation* (Vol. 18, pp. 313-371). New York: Academic Press.
 BYRNE, R. M. J. (1989). Suppressing valid inferences with conditionals. *Cognition*, **31**, 61-83.
 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 vs. syntactic approaches to training deductive reasoning. *Cognitive Psychology*, **18**, 293-328.
 CHENG, P. W., & NOVICK, L. (1990). A probabilistic contrast model of causal induction. *Journal of Personality & Social Psychology*, **58**, 545-567.
 CHI, M. T., FELTOVICH, P. J., & GLASER, R. (1981). Categorization and representation of physics problems by experts and novices. *Cognitive Science*, **5**, 121-152.
 EVANS, J. ST. B. T. (1982). *The psychology of deductive reasoning*. London: Routledge & Kegan Paul.
 EVANS, J. ST. B. T. (1989). *Bias in human reasoning: Causes and consequences*. Hove and London, U.K.: Erlbaum.
 EVANS, J. ST. B. T., BARSTON, J. L., & POLLARD, P. (1983). On the conflict between logic and belief in syllogistic reasoning. *Memory & Cognition*, **11**, 295-306.
 FILLENBAUM, S. (1975). If: Some uses. *Psychological Research*, **37**, 245-260.
 FILLENBAUM, S. (1976). Inducements: On phrasing and logic of conditional promises, threats, and warnings. *Psychological Research*, **38**, 231-250.
 FILLENBAUM, S. (1978). How to do some things with IF. In J. W. Cotton & R. L. Klatzky (Eds.), *Semantic factors in cognition* (pp. 169-214). Hillsdale, NJ: Erlbaum.
 GEIS, M., & ZWICKY, A. M. (1971). On invited inferences. *Linguistic Inquiry*, **2**, 561-566.
 GRICE, H. P. (1975). Logic and conversation. In P. Cole & J. L. Morgan (Eds.), *Syntax and semantics* (Vol. 3, pp. 41-58). New York: Academic Press.
 GRIGGS, R. A., & COX, J. R. (1982). The elusive thematic-materials effect: Wason's selection task. *British Journal of Psychology*, **73**, 407-420.
 HENLE, M. (1962). On the relation between logic and thinking. *Psychological Review*, **69**, 366-378.
 JOHNSON-LAIRD, P. N. (1983). *Mental models*. Cambridge, U.K.: Cambridge University Press.
 KAHNEMAN, D., & MILLER, D. T. (1986). Norm theory: Comparing reality to its alternatives. *Psychological Review*, **93**, 136-153.
 KELLEY, H. H. (1972). Causal schemata and the attribution process. 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. 1-26). Morristown, NJ: General Learning Press.
 KELLEY, H. H. (1973). The process of causal attribution. *American Psychologist*, **78**, 107-128.
 KEPPEL, G. (1973). *Design and analysis*. Englewood Cliffs, NJ: Prentice-Hall.
 LARKIN, J. (1981). *The role of problem representation in physics* (Tech. Rep. No. CIP-429). Pittsburgh, PA: Carnegie-Mellon University, Department of Psychology.
 LARKIN, J., McDERMOTT, J., SIMON, D. P., & SIMON, H. A. (1980). Expert and novice performance in solving physics problems. *Science*, **208**, 1335-1342.
 MANKTELOW, K. I., & EVANS, J. ST. B. T. (1979). Facilitation of reasoning by realism: Effect or non-effect? *British Journal of Psychology*, **70**, 477-488.
 MARCUS, S. L., & RIPS, L. J. (1979). Conditional reasoning. *Journal of Verbal Learning & Verbal Behavior*, **18**, 199-233.
 MARKOVITS, H. (1986). Familiarity effects in conditional reasoning. *Journal of Educational Psychology*, **78**, 492-494.

- MEDIN, D. L. (1983). Structural principles in categorization. In T. J. Tighe & B. E. Shepp (Eds.), *Perception, cognition, and development: Interactional analyses* (pp. 203-230). Hillsdale, NJ: Erlbaum.
- MEDIN, D. L., ALTOM, M. W., EDELSON, S. M., & FREKO, D. (1982). Correlated symptoms and simulated medical classification. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, *8*, 37-50.
- MEDIN, D. L., & SCHAFFER, M. M. (1978). Context theory of classification learning. *Psychological Review*, *85*, 207-238.
- NOVICK, L. R. (1990, April). *The influence of problem structure and world knowledge on causal inference*. Paper presented at the Annual Meeting of the Rocky Mountain Psychological Association, Tucson.
- REICH, S. S., & RUTH, P. (1982). Wason's card selection task: Verification, falsification, and matching. *British Journal of Psychology*, *73*, 395-405.
- REVLIN, R., LEIRER, V., YOPP, H., & YOPP, R. (1980). The belief-bias effect in formal reasoning: The influence of knowledge on logic. *Memory & Cognition*, *8*, 584-592.
- RIPS, L. J. (1983). Cognitive processes in propositional reasoning. *Psychological Review*, *90*, 38-71.
- RUMAIN, B., CONNELL, J., & BRAINE, M. D. S. (1983). Conversational comprehension processes are responsible for reasoning fallacies in children as well as adults: IF is not the biconditional. *Developmental Psychology*, *19*, 471-481.
- WASON, P. C., & JOHNSON-LAIRD, P. N. (1972). *Psychology of reasoning: Structure and content*. Cambridge, MA: Harvard University Press.
- WILKINS, M. C. (1928). The effect of changed material on the ability to do formal syllogistic reasoning. *Archives of Psychology* (No. 102).

(Manuscript received July 5, 1990;
revision accepted for publication October 29, 1990.)