Pragmatic versus Syntactic Approaches to Training Deductive Reasoning

PATRICIA W. CHENG

Carnegie–Mellon University

KEITH J. HOLYOAK

University of Michigan

AND

RICHARD E. NISBETT AND LINDSAY M. OLIVER

University of Michigan

Two views have dominated theories of deductive reasoning. One is the view that people reason using syntactic, domain-independent rules of logic, and the other is the view that people use domain-specific knowledge. In contrast with both of these views, we present evidence that people often reason using a type of knowledge structure termed *pragmatic reasoning schemas*. In two experiments, syntactically equivalent forms of conditional rules produced different patterns of performance in Wason's selection task, depending on the type of pragmatic schema evoked. The differences could not be explained by either dominant view. We further tested the syntactic view by manipulating the type of logic training subjects received. If people typically do not use abstract rules analogous to those of standard logic, then training on abstract principles of standard logic alone would have little effect on selection performance, because the subjects would not know how to map such rules onto concrete instances. Training results obtained in both a laboratory and a classroom setting confirmed our hypothesis: Training was effective only when abstract principles were coupled with examples of selection problems, which served to elucidate the mapping between abstract principles and concrete instances. In contrast, a third experiment demonstrated that brief abstract training on a pragmatic reasoning schema had a substantial impact on subjects' reasoning about problems that were interpretable in terms of the schema. The dominance of pragmatic schemas over purely syntactic rules was discussed with respect to the relative utility of both types of rules for solving real-world problems. © 1986 Academic Press, Inc.

How do people reason about problems in everyday life? Two views have dominated theories of deductive reasoning. According to one view people use syntactic, domain-independent rules of logic. Some philosophers (e.g., Goodman, 1965) and psychologists (including Piaget and his followers) who are sympathetic to this view believe that these rules correspond closely to those in standard logic. Substantial evidence shows, however, that typical college students often do not reason in accord with the rules of standard logic (see Evans, 1982, for a review). Partly in response to this evidence, other proponents of the syntactic view have proposed that people use a *natural* logic which consists of a repertory of syntactic rules that people untutored in standard logic naturally use (e.g., Braine, 1978; Braine, Reiser, & Rumain, 1984; Rips, 1983). The empirical adequacy of the natural-logic approach has also been challenged, however (Johnson-Laird, 1983, chap. 2).

Partly in response to the empirical difficulties encountered by the syntactic view, an opposing view has developed that holds that people do not use syntactic rules at all but instead develop much narrower rules tied to particular content domains in which people have actual experience. Such specific rules, or perhaps simple memory of examples and counterexamples, are used to evaluate the validity of propositions (e.g., Griggs & Cox, 1982; Manktelow & Evans, 1979; Reich & Ruth, 1982).

In contrast to both the extreme syntactic view and the specific-experience view, Cheng and Holyoak (1985) proposed that people often reason using *pragmatic reasoning schemas:* clusters of rules that are highly generalized and abstracted but nonetheless defined with respect to classes of goals and types of relationships. An example of a pragmatic reasoning schema is the set of abstracted rules for situations involving "permission," that is, situations in which some action A may be taken only if some precondition B is satisfied. If the semantic aspects of a problem suggest to people that they are dealing with a permission situation, then all of the rules about permissions in general can be called on, including "If action A is to be taken, then precondition B must be satisfied," "Action A is to be taken only if precondition B is satisfied," "If precondition B is not satisfied, then action A must not be taken," and so on. A related example is the schema for situations involving "obligation," that is, situations in which the occurrence of some condition A incurs the necessity of taking some action B. Rules about obligations are similar to but not identical with rules about permissions. For instance, the rule "If condition A occurs, then obligation B arises" implies "If obligation B does not arise, then condition A must not have occurred," but not "Condition A occurs only if obligation B arises."

The rules composing these schemas are the ones that people call on when solving problems of various kinds. Since the rules of some of the schemas lead to the same solution as do the rules of standard logic, people's answers to problems will often appear consistent with those of standard logic. This consistency does not mean that their answers have been produced by the application of syntactic logical rules, since the same people will at other times produce answers that violate those same rules of logic. This occurs when the rules of the schema used lead to conclusions that differ from those that follow from standard logic. By manipulating the semantic features of the problems to evoke different reasoning schemas, one should be able to manipulate whether or not people's answers are correct according to standard logic.

Cheng and Holyoak obtained several empirical results that speak strongly for the existence of reasoning schemas. Two of these findings involve Wason's (1966) selection task. In one version of the Wason task subjects are presented with four cards, which show an "A", a "B", a "4", and a "7". They are informed that the cards have letters on one side and numbers on the other, and are then given the rule, "If a card has an 'A' on one side, then it has a '4' on the other." The task is to indicate all and only those cards that *must* be turned over to determine whether the rule is violated. Interpreting the *if-then* connective as the material conditional in standard logic, the correct answer in this example is to turn over the cards showing "A" and "7". More generally, the rule used in such problems is a conditional statement, *if p then q,* and the relevant cases are p (because if p is the case it must be established that q is also the case) and *not-q* (because if it is not the case that q it must be established that it is also not the case that p).

When the selection problem is presented in an arbitrary form, as in the above example, fewer than 10% of college students typically produce the correct answer. A frequently chosen pattern is "A and 4." One of the errors in such an answer is omission of the card showing "7", indicating a failure to see the equivalence of a conditional statement and its contrapositive (i.e., "If a card does not have a '4' on one side, then it does not have an 'A' on the other"). Other errors include the fallacy of affirming the consequent (which corresponds to insistence on examining "4", which is unnecessary because the rule does not specify anything about the obverse of cards with a "4" on one side), and denying the antecedent (which corresponds to insistence on examining "B", which also is unnecessary because the rule does not specify anything about cards that do not have an "A" on one side).

Cheng and Holyoak pointed out that the schema for "permission" should be particularly useful in performing the selection task because the rules that comprise it lead to the same responses as follow from the material conditional. In one of their experiments subjects were presented with a selection problem based on an abstract description of a permission situ-

ation: "If one is to take action 'A', then one must first satisfy precondition 'P'." Subjects were also given an arbitrary card version that is syntactically identical to the permission problem. About 60% of subjects solved the abstract permission problem correctly, versus only about 20% who correctly solved the card problem. The fact that a purely abstract description of a permission situation produces facilitation supports the schema hypothesis over proposals based on either purely syntactic rules or specific experiences.

In another experiment, Cheng and Holyoak demonstrated that providing an explicit purpose for a rule that would otherwise seem arbitrary could serve to cue the permission schema and hence facilitate performance. For example, the rule "If a passenger's form says 'Entering' on one side, then the other side must include 'cholera,' " was rationalized by explaining that it involved a regulation requiring that travelers show proof of cholera vaccination in order to enter a country. The benefit conveyed by provision of a purpose is inexplicable according to either the specific-experience view or the syntactic view. Since subjects had no experience with the specific content in question and hence no memory for counterexamples to the rule, improvement due to provision of a purpose could not be attributed to processes consistent with the specific-experience view. On the other hand, improvement could not be attributed plausibly to manipulation of the syntactic properties of the problems either, since the added purpose did not affect the logical structure of the problems.

The notion of pragmatic reasoning schemas, represented as sets of inferential rules (Holland, Holyoak, Nisbett, & Thagard, 1986), is essentially a generalization of similar ideas that have been proposed to explain people's causal and inductive reasoning. Kelley (1972, 1973) suggested that people make causal attributions in part by relying on "causal schemas." These are highly generalized, domain-independent, but not purely syntactic rule systems for analyzing causality. People may have, for example, a schema for multiple sufficient causes, where a given effect could be produced solely by operation of a single cause, but a number of different causes could play that role. People's search for causal candidates and their degree of certainty about the effects of a given causal candidate are influenced by the particular causal schema they bring to the analysis of a situation.

Pragmatic reasoning schemas are also related to Johnson-Laird's concept of "mental models." For example, Johnson-Laird (1983, p. 416) discusses a representation of "ownership" in terms of such inferences as, "If I own something then it is permissible for me to use it." In our terms, the concept of ownership is a pragmatic schema based in part on the yet more general permission and obligation schemas. However, the focus of Johnson-Laird's theory of mental models is not on the role of pragmatic

inference rules in reasoning, but rather on a domain-independent scheme for manipulating sets that are represented in terms of tokens for their individual members. The theory has been developed in most detail as an account of reasoning with syllogisms in which the premises specify arbitrary set relations (e.g., "All of the beekeepers are artists"). The theory attributes reasoning errors in this domain largely to limitations on memory capacity. In contrast, the present theory explains errors (as defined by the dictates of standard logic) in terms of the ease of mapping concrete situations into pragmatic schemas, as well as the degree to which the evoked schemas generate inferences that in fact conform to standard logic.

People not only have deductive reasoning schemas, but also inductive reasoning schemas. Nisbett, Krantz, Jepson, and Kunda (1983) and Fong, Krantz, and Nisbett (1986) have proposed that people often reason about events using sets of heuristics or intuitive rules of thumb that are informal equivalents of statistical rules such as the law of large numbers (LLN), the regression principle, and so on. Like causal schemas, these are clusters of inferential rules for predicting and explaining events. They are invoked for reasoning about events that are perceived to be subject to random variations.

Just as Cheng and Holyoak showed that people can solve selection problems if they are given cues that serve to trigger the permission schema, Nisbett et al. (1983) (as well as Jepson, Krantz, & Nisbett, 1983, and Fong et al., 1986) showed that people can produce solutions that are in accord with statistical requirements when the problem contains cues about variability or chance that encourage the use of statistical heuristics. For example, Fong et al. gave subjects a problem about a businesswoman who eats out frequently when she travels to different cities. She returns often to restaurants where she got an excellent meal on her first trip, but is usually disappointed. Most statistically untrained subjects give a causal explanation such as "her hopes were too high" or "the chefs change a lot." But if a cue as to the variability characteristic of restaurant meals is presented, many more subjects give a statistical explanation, such as "maybe it was just by chance that she got such a good meal the first time; there are probably more restaurants that can turn out an occasional excellent meal than restaurants that consistently do that."

Subjects are not only easily cued to use statistical heuristics, but they can also readily be taught to use, and to improve, their statistical heuristics. For example, subjects with no statistical education only rarely give a statistical answer for the uncued version of the restaurant problem above, but increasing amounts of statistical training increase the likelihood of a statistical answer. The great majority of subjects with postgraduate training in statistics gave a statistical answer to the problem.

Even more dramatically, very brief training session in LLN were

helpful in encouraging subjects to apply LLN to problems such as the restaurant problem (Fong et al., 1986). This was true even for completely abstract training, in which the concepts of "sample" and "population" were defined formally and illustrated using problems involving balls in an urn. Fong et al. argued that such abstract training could be effective because subjects already had an informal, approximate grasp of LLN and an ability to apply it to actual events. Thus, they could take immediate advantage of formal improvements to their intuitive understanding. Consistent with this view, Fong et al. found that brief instruction showing subjects how to model example problems in terms of LLN, *without* any additional instruction in the abstract nature of LLN, was also effective in encouraging subjects to use LLN. Fong et al. argued that this could only be true if subjects already had an intuitive version of LLN and could readily generalize from the examples to improvements in this intuitive version. The present investigations are modeled on those we have just described. We test several hypotheses.

First, if people have pragmatic reasoning schemas corresponding to the "permission schema," the "obligation schema," and so on, then it should be easy to encourage them to use these schemas by presenting them with semantic cues designed to trigger them. Triggering a schema that produces the same solution to an *if-then* problem as does the material conditional, such as the permission schema, should help subjects to solve it. Conversely, triggering a schema that does not produce the same solution as the material conditional should be less helpful.

Second, unlike abstract training in LLN, which produces a marked improvement in peoples' ability to apply LLN to concrete problems, abstract training in the material conditional should produce little improvement in people's ability to apply the rule system to concrete problems. Since in our view the rule system is not used in natural contexts, people lack the requisite skills to interpret problems in terms of the material conditional, and hence would profit little from instruction in it. Similarly, simply providing subjects with a few example problems solved by the application of the material conditional would likewise produce little benefit, since few subjects would be able to spontaneously induce the abstract rules from the examples. Thus, if people are to be able to use the rule system at all they must be given training both on abstract rules and on how to apply the rule system to specific examples.

Third, abstract training in pragmatic reasoning schemas, like abstract training in LLN, and unlike abstract training in the material conditional, should produce improvement in ability to use the schemas. Since rules for interpreting concrete problems in terms of the abstract rules of reasoning schemas already exist, improvement on the abstract rules should carry over to improved ability to solve concrete problems.

EXPERIMENT 1

Experiment 1 was designed to assess the influence of reasoning schemas on performance on the selection task, as well as to assess the usefulness of different training procedures.

If people interpret *if-then* in terms of reasoning schemas, selection problems that evoke different pragmatic reasoning schemas might lead systematically to different response patterns. We tested this prediction by measuring the effect of three types of selection problems on reasoning performance. The problems described relations that we thought would be suggestive of permission situations, relations with a causal flavor that we thought would invite a converse assumption (i.e., assuming *if q then p* when given *if p then q*), and relations that were purely arbitrary. The pragmatic schema hypothesis leads us to expect that performance would be better for permission problems than for problems with a converse bias or arbitrary problems. Performance on converse-bias problems should be no worse than for arbitrary problems, because the latter do not evoke any useful schema. In addition to the above three types of problems, all of which involve conditional rules, we included problems that explicitly stated biconditional rules.

We also assessed the effect of different training procedures. The design with respect to training was the same as that employed by Fong et al. (1986). Subjects received either abstract rule training in the conditional, training in how to apply the conditional to concrete example problems, both, or neither. Consistent with the view that subjects already possessed an intuitive version of LLN, Fong et al. found that abstract rule training was effective by itself, as was example training. In contrast, we anticipated that since people do not possess rules corresponding to the conditional, neither type of training would be effective by itself. Abstract rule training should not be effective because subjects have no means of interpreting problems so that the rules can be applied. Example training, showing how to apply the conditional, should also not be effective by itself, because subjects have no intuitive version of the conditional to begin with.

Method

Subjects

Eighty students at the University of Michigan were randomly assigned in equal numbers to each of four training conditions: (1) Rule training, (2) Examples training, (3) Rule plus Examples training, and (4) No training. None of the subjects had previously received any formal training in logic.

Training Materials

Rule training. Subjects receiving rule training read a seven-page booklet consisting of an exposition on conditional statements, followed by an inference exercise. These materials

are presented in Appendix A. The exposition consisted of an explanation of the equivalence between a conditional statement and its contrapositive, as well as an explanation of the two common fallacies of affirming the consequent and denying the antecedent. The contrapositive was explained in part by the use of a truth table, in part by Euler diagrams that used concentric circles to show the relations between a conditional statement and its contrapositive, and in part by an illustrative conditional statement. The illustrative statement expressed a realistic causal relation. Similarly, the fallacies were explained in part by diagrams and in part by alternative possible causes related to the illustrative statement.

At the end of rule training, subjects were given an *inference* exercise in which they were to select statement(s) that can be validly inferred from each of three given conditional statements. The statements were all in the form of *if p then q*. The randomly ordered possible inferences were in the following forms: *if not-p then not-q* (invalid), *if not-q then not-p* (valid), and *if q then p* (invalid). Subjects were given immediate feedback on correctness, followed by a brief explanation of the correct answer.

Examples training. Subjects receiving examples training were requested to attempt to solve two selection problems. They were given immediate feedback about their performance. One example was the "department store" problem used by D'Andrade (1982), and the other was a problem in which a catalog of paintings had to be checked to determine whether "all the Cubist paintings are by Picasso." Neither example bore any obvious similarities to the later test problems. The correct answer to each example was explained in terms specific to the particular problem.

Rule plus example training consisted of the materials for the rule condition followed by those for the examples condition. The only further addition was that for these subjects the explanation of the correct answer for each example was couched in terms of the abstract rules they had just learned.

Test Problems

Eight selection problems were used, consisting of two of each of three types of problems involving a conditional rule and two problems involving a biconditional rule. The problems are presented in Table 1. Arbitrary problems (the Wason "card" problem and the "bird" problem) bore little relationship to the prior knowledge subjects were likely to have. Converse-bias problems were more realistic; however, subjects' prior knowledge was expected to encourage assumption of the converse (i.e., *if q then p*). The "washing labels" problem involved a regulation for which causal knowledge would suggest an "if and only if" interpretation, and the "electrical charges" problem involved a causal regularity in which only a single cause was likely to be considered, also leading to an "if and only if" reading. The permission problems ("cholera" and "drinking age") are readily interpreted as permissions, which should not encourage assumption of the converse, but instead should yield the same responses as follow from the conditional of standard logic. Finally, the biconditional problems stated explicitly that the converse of a conditional rule was also true. The content of the biconditional problems was relatively arbitrary. The anticipation was that this arbitrary content would block application of any pragmatic reasoning schema and that performance on these problems would be poor.

Each problem described a brief scenario. Within each scenario was embedded a conditional or biconditional rule, a question asking the subject to determine the correctness of the rule, and a list of the four possible cases (*p, not-p, q,* and *not-q*) from which the subject was to select. These cases were randomly ordered.

In order to provide a second measure of the effectiveness of abstract training, subjects were asked to judge which of a series of transformations of a conditional rule retained the basic meaning of the rule itself. The conditional rules in this *equivalence judgment* task were directly excerpted from the selection problems that preceded them. The six equiva-

TABLE 1
If-Then Rules of Various Types and Corresponding Choices in the Selection Problems (Experiments 1 and 2)

Problem type	Rule	Choices (*p, not-p, q, not-q*)
Arbitrary	If a card has an "A" on one side, then it has a "4" on the other side	A B 4 7
	If a bird on this island has a purple spot underneath each wing, then it builds nests on the ground	Bird A has a purple spot underneath each wing Bird B does not have any purple spots Bird C builds nests on the ground Bird D builds nests in trees
Converse bias	If a washing label has "silk" on one side, then it has "dry clean only" on the other side	Silk Cotton Dry clean only Machine wash in warm water
	If two objects carry like electrical charges, then they will repel each other	Two objects that carry like electrical charges Two objects that carry opposite charges Two repelling objects Two objects that do not repel
Permission	If a passenger wishes to enter the country, then he or she must have had an inoculation against cholera	Entering Transit Inoculated against cholera and hepatitis Inoculated against typhoid
	If a customer is drinking an alcoholic beverage, then he or she must be over twenty-one	Customer A is drinking a beer Customer B is drinking tea Customer C is certainly over 50 Customer D looks less than 18
Biconditional	If a card has a circle on one side, then it has the word "red" on the other, and conversely, if it has the word "red" on one side, then it has a circle on the other	(Picture of a circle) (Picture of a triangle) Red Purple
	If a turtle crosses a road, then the flag by the palace flies, and conversely, if the flag by the palace flies, then a turtle crosses a road	A turtle is crossing a road No turtle is crossing any road The flag by the palace is flying The flag by the palace is not flying

lence judgment problems consisted of conditional rules excerpted from the six corresponding selection problems, each followed by a randomly ordered list of four transformations of each of these rules. The rules were always in the form *if p then q*. The forms used as transformations were a simple syntactic variation, *q if p* (valid); the contrapositive, *if not-q then not-p* (valid); the converse, *if q then p* (invalid); and the inverse, *if not-p then not-q* (invalid). Subjects were to judge which of the transformations retained the basic meaning of the corresponding rules. This test was similar to the inference exercise included in the rule and rule plus examples training conditions. Biconditional rules were not included in this test.

Results and Discussions

Results on the selection task were analyzed using analysis of variance with problem type as a within-subject variable and training condition as a between-subject variable. Two measures of performance were analyzed: Whether a subject made at least one error on a problem (i.e., failed to solve the problem perfectly) and whether he or she made each of the four possible kinds of errors. For conditional problems, the four kinds of errors in the selection task were failing to select *p*, failing to select *not-q*, selecting *q*, and selecting *not-p*. These errors correspond, respectively, to errors on modus ponens, modus tollens, affirming the consequent, and denying the antecedent. For biconditional problems the correct response was to select all four alternatives.

Since the two dependent measures of overall performance on a problem—whether a subject made at least one error and the number of errors—reflected the same pattern of results, statistics for only the former measure are reported as an indicator of overall performance on a problem. Performance on conditional and biconditional problems were analyzed separately. For conditional problems, type of problem and training condition did not interact according to either measure of performance ($F < 1$). Accordingly, results for type of problem and training condition are reported separately, each collapsed over the other variable. Table 2 presents the percentages of subjects (collapsed over training conditions) who made errors of various types on each of the four kinds of selection problems. Table 3 presents the percentage of subjects in each training condition (collapsed over conditional selection problems) who made errors of various types on the same task. Half-widths of 95% confidence intervals for various pairwise differences between means are indicated at the bottom of each table of results.

Effects of Problem Type

Consistent with our reasoning schema hypothesis, performance differed markedly across types of conditional problems, $F(2,152) = 64.7$, $p < .001$, for the percentage of subjects making at least one error. Fewest subjects made at least one error on permission problems ($34 \pm 6\%$), more

TABLE 2
Percentage Errors in Selection Task as a Function of Problem Type (Experiment 1)

Problem type	Type of error				At least one error
	p	$not\text{-}q$	q	$not\text{-}p$	
Conditional					
Arbitrary	15 ± 5	62	65	12	81 ± 6
Converse bias	16 ± 5	42	34	17	66 ± 6
Permission	5 ± 5	24	8	2	34 ± 6
Biconditional	14 ± 7	48	14	54	80 ± 9

Note. Error margins indicated in the table are half-widths of 95% confidence intervals for the respective conditions. Numbers in the three center columns have the same half-widths as the left-most numbers in the corresponding rows. The half-width of a 95% confidence interval for pairwise comparisons between mean frequencies of particular errors on the three types of conditional problems is 7. The corresponding half-width for comparisons between the percentages of subjects who made at least one error on various types of problems is 8.

subjects did so on converse-bias problems ($66 \pm 6\%$), and still more subjects did so on arbitrary problems ($81 \pm 6\%$). Moreover, type of error interacted strongly with type of problem, $F(6,456) = 20.7$, $p < .001$. In particular, frequency of erroneously selecting q was highest for arbitrary problems ($65 \pm 5\%$), next highest for converse-bias problems ($34 \pm 5\%$), and lowest for permission problems ($8 \pm 5\%$), confirming our prediction that permission problems have a lower tendency to invite assumption of the converse than the other two types of problems. Interestingly, the frequency of selecting q was much higher for arbitrary than for converse-bias problems. Because arbitrary problems do not evoke any pragmatic schemas, subjects would have to fall back on whatever syntactic logical rules are available or on nonlogical strategies such as "matching" (Manktelow & Evans, 1979). A matching strategy would produce apparent converse errors for our arbitrary problems. Frequency of failure to select *not-q* similarly depended on type of problem, being highest for arbitrary problems ($62 \pm 5\%$), next highest for converse-bias problems ($42 \pm 5\%$), and lowest for permission problems ($24 \pm 5\%$).

Many fewer subjects ($17 \pm 5\%$) made the other two types of errors (i.e., failure to select *p* or erroneous selection of *not-p*) for all three types of conditional problems. Permission problems again produced fewer errors than either arbitrary or converse-bias problems, which produced comparable percentages of errors for these two cases. It is particularly noteworthy that the permission problems yielded more accurate performance even for the choice of *p*, which corresponds to modus ponens,

perhaps the most plausibly psychological of all the syntactic inference rules in standard logic and in Braine's (1978) natural logic.[1]

The two permission problems can be used to test the domain specificity hypothesis, which claims that only rules with which subjects have prior familiarity will yield good performance. For most subjects the "drinking age" rule was presumably much more familiar than the "cholera inoculation" rule. Although the percentage of subjects making at least one error was marginally lower for the more familiar rule (29 ± 7% versus 39 ± 7%, $p = .05$), even the relatively unfamiliar rule produced a much lower error rate than any converse-bias or arbitrary problem ($p < .002$). Thus subjects were able to reason in accord with standard logic even for a relatively unfamiliar rule if it evoked a permission schema. These results indicate that while specific experiences may play a role in reasoning, they cannot account for the large effect of type of problem on performance.

Performance on biconditional problems was poor: 80 ± 9% of the subjects (averaged over training conditions) made at least one error on such problems. Because all four cases should be selected, the four types of errors for biconditional problems all indicate failure to select. As Table 2 indicates, the most frequent errors were failure to select the two negated cases. Biconditional problems, unlike permission problems, cannot be solved by application of the permission schema. Futhermore, the arbitrary formulation of the particular biconditionals used presumably discouraged application of any pragmatic schema, so that subjects' performance on these problems was poor.

Training Effects

It may be seen in Table 3 that the percentage of subjects who make at least one error on conditional problems varied significantly across training conditions, $F(3,76) = 6.04$, $p < .01$. Neither rule training by itself nor examples training by itself was effective. In contrast, rule training coupled with examples training significantly decreased the frequencies of three types of errors—failure to select p, failure to select not-q, and erroneous selection of q. These effects are reflected in a sub-

---

[1] Problem type did not strongly influence errors in the equivalence judgment task. Over all types, the error rate was 33 ± 3% for failing to select the syntactic variation, 48 ± 3% for failing to select the contrapositive, 6 ± 3% for selecting the converse, and 4% for selecting the inverse. Contrary to our expectations, the frequency of converse errors was low even for the converse-bias problems (10 ± 4%), and the frequency of the contrapositive error was not lower for permission problems (46 ± 4%), than for other types. We have no confident explanation for the differences between the effects of problem type observed for equivalence judgments versus selection task performance.

TABLE 3
Percentage Errors in Selection Task as a Function of Training Condition (Experiment 1)

Training condition	Type of error				At least one error
	p	not-q	q	not-p	
Rules & examples	5 ± 9	27	28	8	39 ± 12
Rules only	14 ± 9	48	33	7	65 ± 12
Examples only	10 ± 9	45	37	12	62 ± 12
Control	18 ± 9	51	44	14	75 ± 12

Note. Error margins indicated in the table are half-widths of 95% confidence intervals for the respective conditions. Numbers in the three center columns have the same half-widths as the left-most numbers in the corresponding rows. The half-width of a 95% confidence interval for pairwise comparisons between mean frequencies of particular errors for the three training groups is 12. The corresponding half-width for comparisons between the percentages of subjects who made at least one error for various training conditions is 17.

stantial decrease in the mean percentage of subjects who made at least one error on a problem, from 75 ± 12% to 39 ± 12%.[2]

The above pattern is quite different from that observed by Fong et al. (1986) for the effects of training on use of LLN. They found both rule training and examples training to be effective in isolation, and the effects of both were approximately additive. The present pattern of results supports our contention that the material conditional is not a rule system that subjects use naturally. Thus rule training is ineffective because subjects have no ability to apply it to concrete problems and examples training is ineffective because subjects have no intuitive grasp of the rule they are

---

[2] For the equivalence judgment task, training interacted strongly with type of error, $F(9,228) = 14.0$, $p < .001$. None of the training procedures had any significant effects on the converse and inverse. However, rule training, either alone or coupled with examples, reduced errors on the contrapositive by large and significant amounts, from 73 ± 11% for the control group, to 19 ± 11% for the subjects who received rule training alone, and 28 ± 11% for subjects who received both rule and examples training. Surprisingly, rule training, either alone or together with examples, greatly *increased* errors on the syntactic variation, from 17 ± 11% for the control group, to 55 ± 11% for subjects who received abstract training alone, and 49 ± 11% for subjects who received both rule and examples training. This negative effect of rule training may have resulted from the fact that the inference exercise included as part of the rule training did not include examples of the syntactic variation. It may be that by alerting subjects to fallacies of which they were previously unaware, rule training instilled a general reluctance to accept unfamiliar transformations, leading to the detrimental effect observed. In any case, the fact that rule training affected performance on both the contrapositive and the syntactic variation indicates that the lack of an effect of rule training alone in the selection task was not due to subjects' simply learning nothing from our training protocol. Rule training by itself did alter subjects' performance on a direct inference task, even though what was learned could not be readily applied in the selection task.

being shown how to apply. Because the confidence intervals for pairwise differences between means are quite wide, however, the null hypothesis that neither rule nor examples training alone yielded any benefit cannot be accepted with confidence. In Experiment 2 we examine the effects of a greatly augmented rule-training procedure, namely, exposure to an entire course in logic.

Training did not affect the percentage of subjects who made at least one error on biconditional problems, $F < 1$. Training and type of error did not interact, $F < 1$. These results indicate that training on conditional reasoning did not transfer to solution of biconditional problems. If subjects had an intuitive appreciation of the material conditional, then it might be expected that any advantage gained by training on the conditional would result in some degree of improved understanding of biconditional problems.

In sum, the semantic content of problems significantly influenced reasoning performance. In particular, permission problems yielded better performance than either converse-bias or arbitrary problems, confirming our reasoning schema hypothesis. Training on the conditional in standard logic, when coupled with examples of conditional selection problems, clearly improved performance on subsequent conditional selection problems, though training effects did not transfer to biconditional problems. Training on either standard logic, by itself, or on examples of selection problems, by itself, failed to produce significant improvements, but firm conclusions cannot be drawn from the latter results.

EXPERIMENT 2

The results of Experiment 1 indicated that training in standard logic, when coupled with training on examples of selection problems, leads to improved performance on subsequent selection problems. In contrast, training on rules of logic without such examples failed to significantly improve performance. This is consistent with our view that the material conditional is not part of people's intuitive reasoning repertoire, and hence they lack any ability to put abstract rule training to use. An obvious possibility, however, is that our experimental "micro-course" on the logic of the conditional was simply too minimal to convey much benefit. To assess this possibility, Experiment 2 examined the impact of a much broader and more prolonged abstract training condition: a one-semester undergraduate course in standard logic.

Method

Course Content

Two introductory logic classes provided subjects for Experiment 2. One class was held at the Ann Arbor campus of the University of Michigan and one at the branch campus at

Dearborn. Both classes involved about 40 h of lectures and covered topics in propositional logic, including modus ponens, modus tollens, affirming the consequent, and denying the antecedent, and the distinction between the conditional and the biconditional. The textbook in one class was *Elementary Logic* (Simco & James, 1982), and in the other it was *Introduction to Logic* (Copi, 1982). In both classes the treatment of the valid and invalid inference patterns was primarily formal. While meaningful conditional sentences were introduced in lectures to illustrate the inference rules and fallacies, emphasis was placed on formal logical analyses (i.e., truth-table analyses and construction of proofs). Neither class provided any exposure to the selection task or other psychological research on deductive reasoning.

Procedure and Subjects

A pretest and post-test were administered to each class. The pretest was given in the first week of class before any discussion of the conditional had taken place; the post-test was given in the final week of the semester. To generate matched test materials, the eight selection problems used in Experiment 1 were divided into two sets of four. Each set consisted of one of each of the four problem types (arbitrary, converse bias, permission, and biconditional). The assignment of the two resulting sets of test materials to the pretest and post-test was counterbalanced across subjects.

The test sessions took place during regular class meetings; however, they were introduced to the students as a psychology experiment rather than as part of the course. In each session subjects were asked simply to complete booklets with four selection problems. No feedback of any kind was provided until after completion of the post-test. Only data from students who completed both the pretest and the post-test were analyzed, so that the effect of logic training could be treated as a within-subjects variable. Across the two classes a total of 53 students completed the study.

Results and Discussion

The pattern of results did not differ across the two classes that provided subjects; accordingly, all results are reported using the combined data from both classes. The percentages of subjects who made errors of various types on conditional selection problems (collapsing over arbitrary, converse-bias, and permission types) on the pretest and post-test are presented in Table 4. No significant improvement was obtained in

TABLE 4
Percentage Errors as a Function of Logic Training for Conditional Selection Problems
(Experiment 2)

	Type of error				At least one error
	p	$not\text{-}q$	q	$not\text{-}p$	
Pretest	20 ± 6	55	41	18	75 ± 5
Post-test	14 ± 6	58	31	12	72 ± 5

Note. Error margins indicated in the table are half-widths of 95% confidence intervals for the respective conditions. Numbers in the three center columns have the same half-widths as the left-most numbers in the corresponding rows. The half-width of a 95% confidence interval for pairwise comparisons between mean frequencies of particular errors on the pretest and the post-test is 8. The corresponding half-width for comparisons between the percentages of subjects who made at least one error on the pretest and the post-test is 7.

percentage of problems on which at least one error was made. The mean improvement was a bare $3 \pm 7\%$, $t < 1$. T tests performed on each error type separately revealed that the only type of error for which the improvement was individually significant was erroneous selection of the q alternative, an error that decreased by $10 \pm 8\%$ from pretest to post-test.

No interactions involving training and type of conditional problem approached significance. Performance on biconditional problems showed no sign of improvement between the pretest and post-test. The percentage of subjects making at least one error on the biconditional problem was $81 \pm 8\%$ on the pretest and $87 \pm 8\%$ on the post-test, a nonsignificant *increase* in error rate, $t < 1$. No significant changes in the frequency of individual error types were obtained for the biconditional problems. Overall, then, the only apparent influence of a one-semester logic course was a small decrease in the tendency to make the error corresponding to affirming the consequent (i.e., selecting the q alternative) for conditional selection problems.

The weak effect of a semester's training in logic contrasts dramatically with the effects of semantic variations in description of the selection problem. Table 5 presents the percentages of subjects who made the various types of errors for each of the four problem types, collapsing over the pretest and post-test (since no interactions between logic training and problem type approached significance). The impact of problem type was very similar to that observed in Experiment 1. The three types of conditional problems differed in both percentage of subjects making at least one error and percentage of errors on individual

TABLE 5
Percentage Errors as a Function of Problem Type (Experiment 2)

Problem type	Type of error				At least one error
	p	not-q	q	not-p	
Conditional					
Arbitrary	23 ± 7	71	58	21	91 ± 6
Converse bias	19 ± 7	69	37	19	87 ± 6
Permission	10 ± 7	30	12	6	40 ± 6
Biconditional	24 ± 8	64	25	61	84 ± 9

Note. Error margins indicated in the table are half-widths of 95% confidence intervals for the respective conditions. Numbers in the three center columns have the same half-widths as the left-most numbers in the corresponding rows. The half-width of a 95% confidence interval for pairwise comparisons between mean frequencies of particular errors on the three types of conditional problems is 10. The corresponding half-width for comparisons between the percentages of subjects who made at least one error on various types of problems is 9.

alternatives, $F(2,104) = 63.1$, $p < .001$, for the latter analysis. Performance was much more accurate for the permission type than for the other two conditional types. The interaction between conditional problem type and type of error was also significant, $F(6,312) = 6.90$, $p < .001$. As the data in Table 5 indicate, fewer errors of all types were made for permission problems than for the other problem types. The patterns of errors for the arbitrary and converse-bias types were similar except that fewer erroneous selections of the q alternative were made for the latter type. For all three conditional problems, the most common error was failure to select the *not-q* case.

The pattern of errors for biconditional problems was also similar to that observed in Experiment 1. Most subjects ($84 \pm 9\%$) made at least one error per biconditional problem, and the bulk of the individual errors involved failures to select the two negative cases, *not-q* and *not-p*.

EXPERIMENT 3

We have found that abstract training in the logic of the conditional does not by itself have much effect on the way people reason about problems that could potentially be solved by its use. In contrast, problems that lend themselves to interpretation in terms of pragmatic reasoning schemas are solved by a large fraction of subjects. If this is because people normally solve problems using such schemas, and if Fong et al. (1986) are correct in their assertion that even purely abstract training in naturally occurring rule systems can be effective in encouraging people to use them, then it ought to be possible to improve people's deductive reasoning by training on pragmatic reasoning schemas.

In Experiment 3 we tested this possibility. We provided subjects with an abstract statement of the notion of "obligations," together with a description of the procedures necessary to check whether an obligation has been carried out. The anticipation was that this training would facilitate subjects' solution of problems that are in fact semantically interpretable in terms of the obligation notion. We predicted that subjects would not need training on how to map the abstract rules of obligations onto examples of particular problems. Since, as in the case of the schema for LLN, some form of an obligation schema is naturally induced, mapping rules must already exist. Schema training could nonetheless improve performance in at least three ways: (1) by providing subjects with more general mapping rules for interpreting situations in terms of the obligation schema; (2) by ensuring that the inferential rules attached to the schema are valid; and (3) by simply providing checking procedures consistent with the conditional that can be applied to even arbitrary problems.

It was important to establish that it is not merely the latter checking procedure training which is effective, however. Griggs (1983) has argued

that facilitation can be obtained merely by virtue of orienting subjects toward detecting violations of a rule. Thus, in a separate, "contingency" training condition, we trained subjects in the use of checking procedures for so-called contingencies involving the relation between one event or its absence and another event or its absence. It was anticipated that the latter procedure would have little effect on subjects' solution of semantically interpretable problems, since such problems would be understood either in terms of a reasoning schema that maps onto the conditional, in which case the checking procedures would be redundant, or in terms of a reasoning schema that does not map onto the conditional, but which would entail its own checking procedures that would override any purely arbitrary procedures.

On the other hand, we anticipated that it might be useful for subjects' solutions of arbitrary problems to teach them correct checking procedures in the abstract. So long as subjects are able to map the events in a problem onto the checking procedures, their mechanical application would result in correct solution. Thus, training in an obligation schema was expected to improve both problems that are normally interpreted in terms of such reasoning schemas and problems that are not (because the checking-procedure drill contained in schema training should be effective for arbitrary problems as well as for semantically interpretable ones.) In contrast, the simple checking-procedure training of the "contingency" condition was expected to be useful, if at all, primarily for arbitrary problems.

Method

Subjects and Procedures

Seventy-two University of Michigan undergraduates served as paid subjects. Subjects were randomly assigned in equal numbers to one of three conditions: a control condition in which no training was given, the obligation-training condition, and the contingency-training condition. Subjects in the two training conditions were given the appropriate training materials to read for 10 min and were then asked to complete the test problems. The control subjects were simply given the test problems and asked to complete them.

Training Materials

Obligation training. Subjects in the obligation-training condition received a two-page booklet detailing the nature of obligations and the procedures neccessary for checking if a violation of the obligation has occurred. An example of an obligation statement presented in the *if-then* conditional form was given. The procedures for assessing obligations were described in terms of four rules, one for each of the four possible situations that might arise (situations that can be mapped onto p, not-p, q, and not-q). The full training materials are presented in Appendix B.

Contingency training. The contingency-training materials, also presented in Appendix B, were closely matched to the obligation-training materials, except that the checking proce-

dures were described in terms of assessment of "contingencies" rather than "obligations." Two examples were provided, one of which was the same as that used in the obligation-training materials, and one of which described a contingency between properties associated with a category. The contingency-training materials thus provided one more example than did the obligation-training materials. Checking procedures were again described in terms of four rules that mapped onto the conditional.

Test Problems

The test-problem booklet consisted of eight selection problems, presented in Table 6. Four of the problems are readily interpretable as obligation situations, and four are relatively arbitrary. Subjects always received the obligation problems as a block and the arbitrary problems as a block. Order of the blocks was counterbalanced across subjects within each training condition.

Results and Discussion

Table 7 presents the percentage errors for the two problem types as a function of training condition. These results are based on the percentage of subjects making at least one error on a problem (i.e., failing to give the correct choice, p and not-q). An analysis of variance was performed with the three training conditions as a between-subject factor and the two problem types (artibrary and obligation) as a within-subject factor. Both main effects were highly significant. Subjects made significantly fewer errors on obligation problems than on arbitrary problems, $F(1,69) = 63.0$, $p < .001$. Thus, conditional problems that trigger the obligation schema, like those that trigger the permission schema, result in higher solution rates. Training effects were also highly significant. As anticipated, obligation training produced the best performance, followed by contingency training, followed by the control condition that received no training, $F(2,69) = 5.71$, $p < .005$. Analyses of individual types of errors yielded a similar pattern, with both the error of selecting q and of omitting not-q being decreased by obligation training and by obligation content.

The trends apparent in Table 7 indicate that the benefit conveyed by obligation training was in part specific to problems interpretable as obligations. Individual t tests revealed that for obligation problems, training in obligations produced performance superior either to contingency training or to the no-training control condition, and, as anticipated, the mere checking procedure instruction of the contingency-training condition produced no detectable improvement for these semantically meaningful problems. For the arbitrary problems, on the other hand, the effects of obligation and contingency training did not differ from each other (although only the former produced significantly fewer errors than the control condition). A more fine-grained analysis of the response patterns for the four arbitrary problems provided further evidence that the effect of obligation training varied across problems. For two of the problems

TABLE 6
If-Then Rules of Various Types and Corresponding Choices in the Selection Problems
(Experiment 3)

Problem type	Rule	Choices (p, not-p, q, not-q)
Arbitrary	If a card has an "A" on one side, then it has a "4" on the other side	A B 4 7
	If a bird on this island has a purple spot underneath each wing, then it builds nests on the ground	Bird A has a purple spot underneath each wing Bird B does not have any purple spots Bird C builds nests on the ground Bird D builds nests in trees
	If a bolt of cloth has any red threads in it, then it must be stamped with a triangle	Bolt A has red threads in it Bolt B has no red threads in it Bolt C is stamped with a triangle Bolt D is not stamped with a triangle
	If a house was built before 1979, then it has a fireplace	House A was built before 1979 House B was built after 1979 House C has a fireplace House D does not have a fireplace
Obligation	If a steel support is intended for the roof, then it must be rustproof	Support A is intended for the roof Support B is intended for the foundation Support C is rustproof Support D is not rustproof
	If any urithium miner gets lung cancer, then the company will pay the miner a sickness pension	Urithium miner A has lung cancer Urithium miner B does not have lung cancer Urithium miner C is receiving a company sickness pension Urithium miner D is not receiving a company sickness pension
	If any of you wins an athletic award, then that person will have to treat the others to a round of drinks at Sam's	Person A won a letter in basketball Person B never played any sport Person C treated everyone to a round of drinks several times Person D has never bought a round of drinks
	If one works for the Armed Forces, then one must vote in the elections	Person A works for the Armed Forces Person B does not work for the Armed Forces Person C voted Person D did not vote

TABLE 7
Percentage Errors as a Function of Problem Type and Training Condition (Experiment 3)

	Training condition			
Problem type	Control	Contingency	Obligation	$\overline{X}$
Arbitrary	73	55	45	58 ± 6
Obligation	36	34	8	26 ± 6
$\overline{X}$	54 ± 12	45 ± 12	27 ± 12	

Note. Error margins indicated in the table are half-widths of 95% confidence intervals for the row and column means. The half-width of a 95% confidence interval for pairwise comparisons between row means is 8. The corresponding half-width for pairwise comparisons between column means is 17.

("red threads" and "fireplace"; see Table 6) the average percentages of errors were 60, 50, and 27 for the control, contingency-, and obligation-training conditions, respectively. This pattern resembles that for the obligation problems in that only the difference between obligation training and the other two conditions was significant. It seems possible that some subjects might have been able to interpret these two problems as obligation situations, once they had received training in obligations. In contrast, for the two remaining arbitrary problems ("card" and "bird"), which were the most difficult of all, obligation training reduced errors no more than did contingency training. Percentages of errors were 85, 60, and 58 for control, contingency, and obligation conditions, respectively.

These results contrast in a very clear way with those of the effects of purely abstract training observed in Experiments 1 and 2, in which we found that neither a brief rule-training session in the logic of the conditional nor an entire course in logic had any substantial effect on the way subjects dealt with selection problems. In contrast, the two brief procedures used in Experiment 3 both had substantial effects. First, and most importantly, obligation schema training improved subjects' performance on problems that were readily interpretable in terms of the appropriate obligation schema. The selective benefit of obligation training extended to two of the relatively arbitrary problems which were interpretable in terms of the obligation schema. These results support our reasoning schema hypothesis. Training in a pragmatic reasoning schema encourages use of that schema for problems that compel a semantic interpretation consistent with the schema and, in addition, may refine subjects' understanding of situations that are potentially interpretable in terms of the schema.

Second, both schema-based training and training in a simple checking procedure improved subjects' performance on the most arbitrary problems. Unlike other studies in which orientation toward violation

checking provided no benefit for arbitrary problems (Griggs, 1984; Yachanin, 1985), the present training procedure provided subjects with information as to which cases in fact constituted violations as well as simply orienting them toward checking violations. It is noteworthy that this brief procedure was effective for at least the arbitrary problems, whereas the formal-logic instruction in Experiments 1 and 2 was not.

GENERAL DISCUSSION

Summary

The present results provide support for the view that people typically reason using abstract knowledge structures organized pragmatically, rather than in terms of purely syntactic rules of the sort that comprise standard logic. Subjects reasoned in closer accord with standard logic when thinking about problems intended to evoke permission or obligation schemas than when thinking about purely arbitrary relations. This pattern is rather ironic, since standard propositional logic does not capture the deontic content of permissions and obligations, whereas it could represent relations of the sort described in our arbitrary problems. These results on problem types are incompatible with the domain specificity view because experience with the precise rules referred to in the permission and obligation problems was not necessary for successful performance (Experiments 1 and 3). The results are also incompatible with the syntactic view, because all problem types were stated in syntactically equivalent forms.

Our training results provide further evidence favoring the pragmatic over the syntactic view. An entire course in standard logic had no effect on the avoidance of any error save a slight reduction in the fallacy of affirming the consequent. A brief training session in formal logic, of a type shown to produce substantial effects on people's ability to reason in accord with the law of large numbers (Fong et al., 1986), had no significant effect on subjects' ability to use modus ponens or modus tollens or to avoid the errors of affirming the consequent or denying the antecedent. This was not simply because the rule training was inherently useless, since, when it was combined with examples training, subjects were able to make substantial use of it.

The near total ineffectiveness of purely abstract training in logic contrasts dramatically with the ready ease with which people seem able to apply a naturally acquired pragmatic reasoning schema. After one semester's training in standard logic, the students in Experiment 2 solved only 11% of the arbitrary problems correctly, whereas the same students, prior to receiving any formal training, solved 62% of the permission problems correctly. Moreover, subjects who received a brief training

session on the obligation schema improved markedly on selection problems interpretable in terms of that schema.

The generality of the benefit apparently conveyed by evocation of a pragmatic schema is also striking. The permission problems yielded significantly fewer errors of all types, including not only the common error of failing to select *not-q* (equivalent to modus tollens), but also the much less frequent error of failing to select *p* (equivalent to modus ponens). In contrast to the benefit conveyed by the evocation of a permission schema, a course in logic produced no significant reduction in either of these errors. The failure to reduce the frequency of errors for modus ponens cannot be attributed to a floor effect, since evocation of the permission schema did reduce the frequency of errors for the *p* alternative. This failure of abstract training to facilitate use of modus ponens suggests that even this rule may not be a general rule of logic for at least some subjects. The evidence that modus ponens can be overridden by a matching strategy (Manktelow & Evans, 1979; Reich & Ruth, 1982) also supports this hypothesis. If modus ponens is not a universal rule of natural logic, as our results suggest, it is difficult to imagine any formal deductive rule that is general across the adult population. In fact, some logicians (e.g., Anderson & Belnap, 1975; Nelson, 1933) have objected to indiscriminate use of modus ponens to draw inferences across irrelevant propositions. Since pragmatic deductive schemas place semantic restrictions on the antecedent and the consequent, inferences involving irrelevant propositions do not occur.

Possible Criticisms of the Schema Hypothesis

A critic might argue that to explain performance on the selection task, one could simply consider the biconditional or converse bias of various *if-then* statements. (Biconditional bias refers to the tendency to invoke a *p if and only if q* interpretation of a given statement *if p then q;* converse bias refers to the tendency to invoke the *if q then p* assumption without entailing the contrapositive of either the given *if p then q* statement or the assumed *if q then p* statement.) Both permission and obligation problems tend not to invite the biconditional or converse assumption, and one might thus argue that their conditional bias is sufficient to explain facilitation. Henle (1962) has suggested that if invited assumptions are taken into consideration, adults if fact reason in accord with standard logic. Others have pointed out that certain contexts invite the biconditional interpretation of *if-then*, whereas other contexts do not (Fillenbaum, 1975, 1976; Geis & Zwicky, 1971). What, then, warrants an emphasis on the notion of reasoning schemas?

First, an account in terms of invitation of the biconditional cannot in fact explain selection performance. As many investigators have pointed

out, most patterns of selection are irreconcilable with either the conditional or the biconditional interpretation of *if-then*. Moreover, our own results show that error rate is high even when the problem in fact requires treatment as a biconditional (Experiments 1 and 2). Thus, no interpretation in terms of standard logic can explain our results. Neither can an interpretation in terms of any natural logic. In particular, no natural-logic interpretation can explain why the *not-q* case was selected more often for permission and obligation problems than for converse-bias and arbitrary problems. Decision on the *not-q* case is not logically dependent on the converse or biconditional assumption, and although a natural logic such as Braine's (1978; Braine et al., 1984) can explain omission of the *not-q* case by some subjects, it cannot explain both the inclusion of the *not-q* case for some problems and its omission in other problems by the same subjects.

Second, *if-then* statements that have a conditional bias differ in regard to the appropriateness of inferring *p only if q* from *if p then q*. Cheng and Holyoak (1985) showed that when subjects were asked to rephrase *if-then* statements into a logically equivalent form, permission statements such as "If you are to enter this country, then you must have an inoculation against cholera" were often rephrased into "You are to enter this country only if you have an inoculation against cholera." Now consider an obligation statement such as, "If an employee is injured at work, then the company must pay for medical expenses." Applying the same syntactic transformation as in the above permission statement, we get "An employee is injured at work only if the company pays for medical expenses." This rephrasing sounds anomalous, because it suggests that the company's payment is a precondition for an employee getting injured. More generally, from an obligation statement of the form, "If situation *P* obtains, then one must fulfill obligation *Q*," one would not want to infer, "Situation *P* obtains only if one fulfills obligation *Q*." In contrast, from a permission statement of the form "If one is to take action *P*, then one must satisfy precondition *Q*," one can quite appropriately infer, "One is to take action *P* only if one satisfies precondition *Q*." Thus, an inference from *if p then q* into *p only if q* depends crucially on the reasoning schema onto which an *if-then* statement is mapped, even among statements that have a conditional bias (also see Wason & Johnson-Laird, 1972, pp.73–75). This state of affairs is inexplicable in terms of either standard logic or natural-logic models.

Another potential criticism of our schema hypothesis might be based on the fact that the schemas that produced the most facilitation in our experiments (permission and obligation) all involved checking for violations of established rules, rather than testing whether rules were true or false. Is there any need to invoke the schema notion rather than simply to suppose that violation checking is easier than hypothesis testing?

We believe the latter alternative is inadequate for the following reasons. As we pointed out in discussing the results of Experiment 3, it appears that orientation toward violation checking does not improve performance unless it is made clear what cases would in fact constitute violations. This, of course, is information that regulation schemas of the sort we have discussed are able to provide. Moreover, obligation training was more effective than contingency training, even though both procedures were concerned with violation checking. Therefore, improved violation checking cannot account for our results. In addition, the schema hypothesis predicts that if rules are mapped onto a schema that corresponds closely to standard logic, performance in hypothesis testing as well as violation checking can be improved. In particular, if a rule is mapped onto a causal schema that suggests the presence of multiple deterministic causes, people should select relevant evidence in accord with the dictates of standard logic (i.e., judging the *q* case to be irrelevant to the truth of the rule, and the *not-q* case to be potentially falsifying). We would predict, for example, that such a pattern would be observed in evaluating a hypothesis such as, "If Peter has thrown a vase at Paul, then the vase is broken." In contrast, if the rules in an evoked schema do not correspond closely to those in standard logic, then little facilitation will result even if the rules check for violations (see Cheng & Holyoak, 1985).

Relations between Pragmatic Schemas and Other Modes of Deductive Reasoning

Our results speak strongly for the existence of pragmatic schemas, since the findings are inexplicable in terms of either the domain specificity view or the domain independence view. It is nonetheless conceivable that multiple types of knowledge relevant to deductive reasoning coexist within a population and even within an individual. Our results do not rule out this interpretation. First, as in other reasoning studies, most of our subjects were correct on modus ponens, even in arbitrary problems, whereas very few were correct on modus tollens. Although modus ponens may not be universal, as discussed above, our results do not exclude the possibility that *some* people may in fact reason with this syntactic rule. (See Rips & Conrad, 1983, for evidence of individual differences in the use of deductive rules.) Moreover, the same individuals who use ponens as a syntactic rule may use tollens only within the context of certain pragmatic schemas. Second, familiarity with a rule may facilitate performance to some extent, as suggested by the marginal difference in selection performance between the two permission problems in Experiment 1. The presumably more familiar "drinking age" rule yielded slightly better performance than the "cholera inoculation" rule. Since familiarity with counterexamples has been shown to be insufficient

by itself for facilitation (Manktelow & Evans, 1979), the effect of remembered counterexamples must be indirect. It seems that familiarity might produce facilitation by ensuring that a counterexample was experienced as such in the context of a reasoning schema, so that when it is remembered, the schema is indirectly evoked.

If pragmatic schemas and syntactic rules coexist, within a population as well as within an individual, what determines when alternative types of knowledge will be used? We propose that if a rule is interpretable in terms of a pragmatic reasoning schema, then a subject will apply the schema. If the rule is not interpretable in terms of a reasoning schema, however, the subject might fall back on whatever syntactic rules are available, or on nonlogical strategies such as "matching".

Why Formal Deductive Rules are Difficult to Induce

In view of our negative conclusion regarding the prevalence of a natural logic based on syntactic rules, an obvious question arises: Why are such rules apparently difficult to induce? Many logicians throughout the centuries have assumed the existence of a natural repertoire of purely syntactic logical rules, as have psychologists such as Piaget. Yet apparently Piaget may have been wrong about people having formal operations of deductive logic, but right about them having a schema corresponding to the inductive rule system embodied in the law of large numbers.

The difficulty of inducing deductive rules may appear paradoxical, since one typically thinks of deductive rules as being more trustworthy than inductive rules. The difficulty is probably not due to greater complexity of the deductive rules, since there is no a priori reason to think that a rule such as modus tollens is more complex than a rough version of LLN. Moreover, people are able to apply the equivalent of modus tollens in some contexts such as permissions.

We believe that the contrast between induction of an intuitive version of LLN and failure to induce the material conditional reflects the relatively narrow range of applicability of the material conditional. A rule such as, "The information from a random sample is likely to be more reliable when the sample is large than when it is small," is virtually context independent. The sample can consist of a collection of any stochastic events specifying any kind of relationship in any domain. In contrast, *reliable, useful* constancies in deductive rules do not hold at a context-independent level. In particular, the abstract concept of the material conditional—i.e., the formal type of contingency situation invented by logicians—does not have universal pragmatic value. It has been pointed out that the "fallacies" of denying the antecedent and affirming the consequent often lead to pragmatically useful inferences in many contexts (Filenbaum, 1975, 1976; Geis & Zwicky, 1971). Moreover, the "valid" inferences are not applicable or useful under some conditions. An obligation

statement in the form *if p then q* seems anomalous when transformed into *p only if q*, a transformation based on a rule in standard logic. More generally, statements that can correspond to the symbols *p* and *q* in the rule *if p then q* are restricted to those that can be true or false (e.g., Quine, 1974). Imperative and interrogative statements are obviously ruled out. (For example, it makes no sense to apply the contrapositive transformation to the sentence, "If John cooks dinner, will Susan do the dishes?") For the same reason, rules associated with the material conditional do not apply to counterfactual, probabilistic, future contingent, or deontic statements. For example, since a deontic statement such as, "If a burglar breaks into your house, then you should call the police," is right or wrong rather than true or false, it is outside the scope of standard logic (although extended logics may describe such cases).

Application of rules associated with the material conditional to statements that are not truth functional may lead to anomalous or useless inferences. For example, the contrapositive transformation of the above deontic statement gives us: "If it is not the case that you should call the police, then a burglar does not break into your house"—an inference that is useless at best. Similarly, the contrapositive transformation of the future contingent sentence, "If the bomb explodes, then everyone will die," gives us, "If not everyone will die, then the bomb does not explode." Whereas the original sentence, although not truth functional, is a meaningful causal prediction, the contrapositive reverses the cause and the effect, yielding an anomalous statement (Cheng & Holyoak, 1985).

Not only are rules associated with the material conditional restricted in their range of applicability, they are never applied in a natural context to antecedents and consequents that are irrelevant to each other, as the syntactic form of the rules ought to allow. For example, according to standard logic the following statements are all true: (1) If France is in Europe, then the sea is salt; (2) If France is in Australia, then the sea is salt; (3) If France is in Australia, then the sea is sweet (Quine, 1974, p. 21). However, even though the formulation of the material conditional was motivated by a desire to represent mathematics in a rigorous way, it is not applied irrelevantly even in the realm of mathematics.[3]

---

[3] Anderson and Belnap (1975, pp. 17–18) relate the following hypothetical anecdote: "A mathematician writes a paper on Banach spaces, and . . . concludes with a conjecture. As a footnote to the conjecture, he writes: 'In addition to its intrinsic interest, this conjecture has connections with other parts of mathematics which might not immediately occur to the reader. For example, if the conjecture is true, then the first order functional calculus is complete; whereas if it is false, then it implies that Fermat's last conjecture is correct.' . . . The editor counters . . . 'In spite of what most logicians say about us, the standards maintained by this journal require that the antecedent of an 'if . . . then' statement must be *relevant* to the conclusion drawn.' " Anderson and Belnap (1975) have argued that relevance is critical to one's intuitive sense of what it is for an argument to be valid.

Given the above restrictions on applicability, it should not be surprising if people typically do not induce some of the syntactic deductive rules associated with the material conditional. Unlike inductive rules such as the LLN, deductive rules are not context independent. But since deductive rules such as those associated with the material conditional do hold some of the time, one might still ask, do people induce syntactic "probabilistic deductive rules" (a concept that might seem inherently contradictory), just as they apparently induce some syntactic probabilistic inductive rules? Our results do not rule out the possibility, any more than they rule out the possibility that people may have some "deterministic" syntactic deductive rules; however, they suggest that regardless of the exact form of the syntactic rules induced, such rules are overridden most of the time by reasoning schemas that pragmatically separate conditions under which the rules hold from conditions under which they do not.

Implications for Education

Our results have clear educational implications. We have shown that deductive reasoning is not likely to be improved by training on standard logic. It seems that since there are many conditions under which the formal rules do not apply, clarifying when the rules do and do not apply is important if students are not to be confused. Unfortunately, the concept of truth functionality seems to be extremely difficult to teach, as a small sample of logicians we asked concurred. And, although the material conditional is useful under certain conditions, people seem to lack the knowledge required to map particular situations onto the syntactic rules. The only comfort given by the present results to the possibility that the conditional may be trainable, in a way that renders it useful for some novel problems, is that in Experiment 1 we found that when examples training was combined with abstract training, there was some improvement on the selection task. The most obvious implication of this result is that, if logic instructors wish to influence their students' inferential behavior in the face of novel problems, they must do much more than they currently do to show how to apply logical rules to concrete problems.

But in our view, the material conditional is a largely artificial reasoning tool that does not capture and purify natural reasoning. Rather, it offers an alternative to it that has pragmatic utility perhaps only for very specialized problems of a kind that do not occur frequently in everyday life. As an approach to improving everyday reasoning, training on pragmatic reasoning schemas seems to us to be more promising. An advantage of training based on naturally induced reasoning schemas lies in the prior existence of interpretative rules for mapping specific situations onto schemas. Moreover, the schemas specify the conditions under which certain clusters of rules would apply. Education in reasoning is most

likely to be effective when it serves to refine pragmatically useful rules that most people will have naturally induced in at least a rudimentary form from everyday experiences.

APPENDIX A: TRAINING MATERIALS USED IN EXPERIMENT 1

Abstract and Abstract Plus Examples Conditions

In this study we are interested in how people interpret and reason about a very important type of logical statement, called the *conditional*. Even though conditional statements are really very simple, people often make errors in dealing with them. These instructions are intended to help you understand conditional statements. Read through these instructions carefully; they should help you solve some reasoning problems you will receive afterward.

A conditional statement consists of two component statements which are often joined by the conjunctions "if . . . then." The conditional statement can be expressed in the standard form.

$$\text{If } p, \text{ then } q \qquad \text{(a)}$$

where it is understood that the letters "p" and "q" each represent a statement. Statement (a) means "if statement p is true, then statement q is also true." For example, let p stand for "It is raining," and let q stand for "The pavement is wet." Then (a) says "If it is raining, then the pavement is wet."

Statement (a) can be expressed in a variety of ways. We will use a horizontal bar above a letter to indicate that a statement is not true. For example, "$\bar{p}$" means "not p." One way of reformulating (a), then, is

$$\text{If } \bar{q}, \text{ then } \bar{p}. \qquad \text{(b)}$$

This means "if statement q is false, then statement p is also false." Rephrasing the above example into form (b) gives, "If the pavement is *not* wet, then it is *not* raining."

People often don't realize at first that statements (a) and (b) are equivalent (identical to each other). To understand the equivalence of statements (a) and (b), consider the circumstances under which (a) is true. The truth of "*If p, then q*" depends on the truth of p and q. The table below lists the truth values of various statements. Reading across and down the table, we see that when p is true and q is also true, then (a) is true (first line). When p is true and q is false, then (a) is false (second line), since (a) says that p implies q. So, in order for (a) to be true when q is false, p cannot be true (comparing the second and third lines). In other words, (a) implies "*If $\bar{q}$, then $\bar{p}$.*"

p	q	(a): *If p, then q*
T	T	T
T	F	F
F	F	T

Another way of understanding the equivalence of (a) and (b) is through the concept of *sets*. In the following diagrams,[4] we will use circles to represent a set, or class of things. We will put circles within circles to represent the inclusion of one class within another. For example, consider Fig. 1 below:

---

[4] The materials actually presented are omitted to conserve space.

[Euler diagram inserted here]

Figure 1 represents the fact that oranges are a subset of citrus fruit, which in turn are a subset of things containing vitamin C.

Now, let p stand for the statement "x is an element in the set P," and let q stand for the statement "x is an element in the set Q." That is, "If p, then q" means "If x is an element in the set P, then x is also an element in the set Q." For what set relation would this statement be true? If p is a subset of Q, as illustrated in Fig. 2a, then every element in P would be an element in Q.

[Euler diagrams inserted here.]

From the same inclusion relation, we can see that if an element x is *not* in the set Q, it cannot be in the set P (see Fig. 2b). Thus, again we see that (a) is equivalent to (b) (i.e., "If p, then q" is equivalent to "If $\overline{q}$, then $\overline{p}$").

To check your understanding of the conditional statement, please answer the question below.

Statement (a), "If p, then q," can be rephrased without changing its basic meaning. Which of the following is a correct rephrasing of (a)? Put a check next to the correct rephrasing(s) before checking the answer on the next page.

( ) 1. *If q, then p.*
( ) 2. *If $\overline{p}$, then $\overline{q}$.*
( ) 3. *If $\overline{q}$, then $\overline{p}$.*

Only 3 is a correct rephrasing. You should note that statement (a) "If p, then q," does *not* imply

$$If\ q,\ then\ p. \qquad (c)$$

It is a common error to assume that (a) implies (c). Rephrasing the example on Page 1 into form (c) gives "If the pavement is wet, then it is raining," which does *not* logically follow from "If it is raining, then the pavement is wet." The pavement may get wet from lawn sprinklers nearby, for instance. In terms of set relations, statement (c) would be true only if the set P includes the set Q, as in Fig. 3:

[Euler diagram inserted here.]

So it should be clear that "If p, then q" does *not* imply "If q, then p."

It is also a common error to assume that "If p, then q" implies

$$If\ \overline{p},\ then\ \overline{q}. \qquad (d)$$

Rephrasing the example on Page 1 into form (d) gives "If it is *not* raining, then the pavement is *not* wet," which again does not follow from "If it is raining, then the pavement is wet," for the same reason we mentioned earlier (e.g., a lawn sprinkler might have made the pavement wet even though it isn't raining).

To sum up, these are the most important facts you need to know about the conditional statement. First, statement (a) is equivalent to statement (b): "If p, then q" implies "If $\overline{q}$, then $\overline{p}$." Second, statement (a) is *not* equivalent to either statement (c) or statement (d): "If p, then q" does *not* imply either "If q, then p" or "If $\overline{p}$, then $\overline{q}$."

Rephrasing Exercise

This exercise will check your understanding of the conditional statement. Which of the

statement(s) below follow logically from the statement, "If the tablecloth is brown, then the wall is white?" Please put a check next to the correct statement(s) below before checking the answer on the next page.

( ) 1. If the tablecloth is *not* brown, then the wall is *not* white.
( ) 2. If the wall is *not* white, then the tablecloth is *not* brown.
( ) 3. If the wall is white, then the tablecloth is brown.

Only 2 is correct. To see this, we can reformulate the statement into the form "If p, then q" by substituting p for "the tablecloth is brown" and q for "the wall is white." Then we see that 2 is in the form "If $\overline{q}$, then $\overline{p}$," which we saw earlier is equivalent to "If p, then q." But 1 is in the form "If $\overline{p}$, then $\overline{q}$," which does *not* follow from "If p, then q." And 3 is in the form "If q, then p," which likewise does *not* follow from "If p, then q."

Below are two more rephrasing problems. When you are done, check the answer on the next page. Statement: If the cube is plastic, then the sphere is metallic. Rephrasings:

( ) 1. If the sphere is metallic, then the cube is plastic.
( ) 2. If the cube is *not* plastic, then the sphere is *not* metallic.
( ) 3. If the sphere is *not* metallic, then the cube is *not* plastic.

Statement: If the beach is white, then the music is slow. Rephrasings:

( ) 1. If the beach is *not* white, then the music is *not* slow.
( ) 2. If the music is slow, then the beach is white.
( ) 3. If the music is *not* slow, then the beach is *not* white.

Answer to rephrasing problems: Only 3 is correct in the above two problems.

Examples and Abstract plus Examples Conditions

The two examples below illustrate how the conditional statement is used to solve problems.

Example 1

As part of your job as an assistant at Sears, you have the task of going through customers' checks to make sure that any check $30 or over has been approved by the section manager. The amount is written on the front of the check, while the section manager's approval is initialed on the back of the check. Which of the checks below would you have to turn over to make sure that the sales clerk has followed the rule? Turn over only those which you need to check the rule. Mark an × below the check(s) you would have to turn over.

[Alternatives listed here.][4]

Please think carefully and solve the problem before checking the answer on the next page.

Answer to Example 1

Conditional statement: If the check is $30 or over, then it has to be approved by the section manager. The correct answer is (a) and (c). Check (a) obviously needs to be turned over. Check (c) also needs to be turned over, because the rule would be violated if the check is over $30. If the check is less than $30, as in (b) and (e), we do not care whether or not it is approved. And if the check is already approved, as in (d), it doesn't matter whether or not it's over $30.

Next Paragraph, Abstract Plus Examples Condition

To reformulate the rule into the form "If p, then q," we substitute p for "the check is $30 or over," and q for "it has to be approved by the section manager." Since alternative (a) corresponds to p, we clearly have to check whether q follows. Since (c) corresponds to $\overline{q}$,

and "*If $\overline{q}$, then $\overline{p}$*" is equivalent to "*If p, then q,*" we have to check whether $\overline{p}$ follows. But we need not turn over checks (b) and (e), since they correspond to $\overline{p}$, and, as we said earlier, "*If p, then q*" does not imply "*If $\overline{p}$, then $\overline{q}$.*" Nor do we need to turn over check (d), which translated into q. As you learned earlier, "*If p, then q*" is *not* equivalent to "*If q, then p*"; accordingly, the other side of check (d) is irrelevant to the truth of the rule.

Example 2

You are helping to compile illustrations for a book on modern painting. The editor's intention is to illustrate each style with works of several painters, rather than a single painter. She suspects, however, that she might have made a slip. She is certain of having included several Picasso cubist paintings, but is not sure whether she has illustrated cubism by any painter other than Picasso. To check this possibility, she asks you to go through a card catalog which you've kept for the list of illustrations. On one side of each card is written the name of the painter, and on the other side is written the style and title of a painting. Which of the four cards below would you need to turn over in order to check her suspicion that all the cubist paintings are by Picasso? In other words, you need to find out whether it is the case that if the painting is cubist, then it is a Picasso.

[alternatives listed here]

When you are done, check the answer on the next page.

Answer to Exercise 2

Conditional statement: If the painting is cubist, then it is a Picasso. The correct answer is to turn over cards (c), (d), and (e). Card (c) should obviously be checked. Cards (d) and (e) should also be checked, since Metzinger and Leger might be cubist painters, in which case the conditional statement would be violated. Card (a) need *not* be checked since it won't help if the Picasso painting is cubist and it is irrelevant if the Picasso painting is not cubist. The surrealist card is obviously irrelevant.

Next Paragraph, Abstract Plus Examples Condition

To reformulate the above conditional statement into the form "*If p, then q,*" we substitute p for "the painting is cubist," and q for "the painting is a Picasso." Since card (c) translates into p, we clearly have to check whether q follows. And since cards (d) and (e) translate into $\overline{q}$, and "*If $\overline{q}$, then $\overline{p}$*" is equivalent to "*If p, then q,*" we have to check whether $\overline{p}$ follows. But we need not turn over card (a), since it corresponds to q, and "*If q, then p*" does not follow from "*If p, then q.*" Accordingly, the backside of card (a) is irrelevant to the truth of "*If p, then q.*" Similarly, we need not turn over card (b), which corresponds to $\overline{p}$, a condition whose implications are irrelevant to the truth of "*If p, then q.*"

APPENDIX B: TRAINING MATERIALS USED IN EXPERIMENT 3

Obligation Condition

In this experiment you will be asked to solve a series of problems involving decisions about what needs to be checked in various situations in order to see whether specified regularities hold. In particular, you will need to apply what you know about regulations involving *obligations* of various sorts. In order to help you with the problems that will follow, please study the following description of the rules required to properly enforce obligations. Obligations are deceptively simple. Thus, although you probably already know what obligations are, these rules will help you think about such situations.

Obligations

As you know, an obligation arises whenever it is the case that certain circumstances or situations create an obligation to perform some action. Obligations can often be stated in an "*If . . . then*" form. For example, the following regulation specifies an obligation: "If a student is a psychology major, then the student must take an introductory psychology course." More generally, if we call the initial situation *I* and the action *C*, an obligation has the form, "If *I* arises, then *C* must be done." In our first example, *I* is "being a psychology major," and *C* is "taking an introductory psychology course."

In order to assess whether an obligation is being satisfied, we need to consider the four possible situations that might arise. These are

1. *I* occurs.
2. *I* doesn't occur.
3. *C* is done.
4. *C* is not done.

Corresponding to each of these possible situations is a rule related to the fulfillment of the obligation. These rules are the following:

1. If *I* occurs, then it is obligatory to do *C*. Clearly, if *I* arises then failure to take the required action would constitute a violation of the obligation. To use our example, if a student is a psychology major, then that student must take an introductory psychology course.

2. If *I* does not occur, then the obligation does not arise. Consequently, *C* need not be done, although the person may do *C* anyway. For example, if a student is not a psychology major the student is not obliged to take an introductory psychology course. It may be permissible, however, for an English major to take an introductory psychology course. But in any case, the basic obligation is simply irrelevant if the student is not a psychology major.

3. If *C* is done, then the obligation is certainly not violated, regardless of whether or not *I* has occurred. If *I* did occur, then the obligation is satisfied. If *I* didn't occur, then the obligation didn't even arise (Rule 2). For example, if we know a student has taken an introductory psychology course, we can be sure the obligation has not been violated: Either the student was a psychology major, and hence fulfilled the obligation, or the student was not a psychology major, in which case the obligation didn't arise.

4. If *C* has not been done, then *I* must not have occurred. This is because if *I* had occurred, then the failure to do *C* would constitute a violation of the obligation. Thus, if a student has not taken an introductory psychology course, the student must not be a psychology major, or else the obligation will have been violated.

If you understand the above four rules, you should find it very easy to assess whether or not an obligation is being met. Note that there are only two situations in which it is possible for an obligation to be violated: When *I* occurs (and *C* is not done) (Rule 1), and when *C* is not done (and *I* occurs) (Rule 4). In the other two situations the obligation can't be violated. These are the cases in which *I* doesn't occur (in which case the obligation doesn't arise) (Rule 2), and in which *C* is done (in which case the obligation will have been met if it arose) (Rule 3).

You may wish to reread these instructions carefully in order to be sure you understand the rules for evaluating obligations. You will then be able to apply what you learned to the test problems. The test problems will include both obligations and other similar types of regularities. You will find it easy to solve these problems if you carefully apply Rules 1–4.

Contingency Condition

In this experiment you will be asked to solve a series of problems involving decisions about what needs to be checked in various situations in order to see whether specified

regularities hold. In particular, you will need to apply what you know about situations involving *contingencies* of a certain sort. In order to help you with the problems that will follow, please study the following description of the rules required to properly evaluate contingencies. Contingencies are deceptively simple. Thus, although you probably already know what contingencies are, these rules will help you think about such situations more systematically.

Contingencies

A *contingency* arises whenever it is the case that a certain condition implies some necessary consequence. Contingencies can be stated in an "*If . . . then*" form. For example, the following statement specifies a contingency: "If a mushroom is red, then it is edible." Another example would be "If a student is a psychology major, then the student must take an introductory psychology course." More generally, if we call the initial condition *I* and the consequence *C*, a universal contingency has the form, "*If I, then C.*" In our first example, *I* is "red mushroom" and *C* is "edible." In the second example, *I* is "being a psychology major," and *C* is "taking an introductory psychology course."

In order to assess whether a contingency in fact holds, we need to consider the four possible situations that might arise. These are

   1. *I* is obtained.
   2. *I* is not obtained.
   3. *C* is obtained.
   4. *C* is not obtained.

Corresponding to each of these possible situations is a rule related to the truth of the contingency. These rules are the following:

1. If *I* is obtained, then *C* must be obtained. Clearly, if *I* is obtained then for *C* not to obtain would show that the contingency doesn't hold. To use our examples, if a mushroom is red, then it must be edible or else the contingency is false; and if a student is a psychology major, then the student must take an introductory psychology course.

2. If *I* is not obtained, then the contingency is not tested. *C* need not obtain, although it may. For example, if a mushroom is not red, it need not be edible, although it may be (perhaps brown mushrooms are also edible). Similarly, if a student is not a psychology major the student need not take an introductory psychology course. It may be possible, however, for an English major to take an introductory psychology course. But in any case, the basic contingency is simply irrelevant if the student is not a psychology major.

3. If *C* is obtained, then the contingency is certainly not falsified, regardless of whether or not *I* obtains. If *I* did occur, then the contingency is satisfied. If *I* is not obtained, then the contingency wasn't even tested (Rule 2). For example, if we know a certain mushroom is edible, we can be sure the contingency was not falsified regardless of the mushroom's color: Either the mushroom is red, and the contingency is satisfied, or it is not red, in which case the contingency was not tested. Similarly, if we know a student has taken an introductory psychology course, we can be sure the contingency has not been falsified: Either the student was a psychology major, and hence satisfied the contingency, or the student was not a psychology major, in which case the contingency wasn't tested.

4. If *C* is not obtained, then *I* must not obtain or else the contingency is falsified. This is because if *I* had occurred, then the failure to obtain *C* would falsify the contingency. Thus, if a mushroom is not edible, it must not be red or else the contingency is false. And if a student has not taken an introductory psychology course, the student must not be a psychology major, or else again the contingency will be falsified.

If you understand the above four rules, you should find it easy to assess whether or not a contingency is being satisfied. Note that there are only two situations in which it is possible for a contingency to be falsified: When *I* is obtained (and *C* is not obtained) (Rule 1), and

when *C* is not obtained (and *I* is obtained) (Rule 4). In the other two situations the contingency can't be falsified. These are the cases in which *I* does not obtain (in which case the contingency is not tested) (Rule 2), and in which *C* is obtained (in which case the contingency will have been satisfied if it was tested) (Rule 3).

You may wish to reread these instruction carefully in order to be sure you understand the rules for evaluating contingencies. You will then be able to apply what you learned to the test problems, which will include a variety of types of such contingencies. You will find it easy to solve these problems if you carefully apply Rules 1–4.

REFERENCES

Anderson, A. R., & Belnap, N. D., Jr. (1975). *Entailment* (Vol.1). Princeton, NJ: Princeton Univ. Press.

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.

Cheng, P. W., & Holyoak, K. J. (1985) Pragmatic reasoning schemas. *Cognitive Psychology*, 17, 391–416.

Copi, I. M. (1982). *Introduction to logic* (6th ed.). New York: Macmillian Co.

D'Andrade, R. (1982, April). *Reason versus logic*. Paper presented at the Symposium on the Ecology of Cognition: Biological, Cultural, and Historical Perspectives. Greensboro, NC.

Evans, J. St. B. T. (1982). *The psychology of deductive reasoning*. London: Routledge & Kegan Paul.

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.

Fong, G. T., Krantz, D. H., & Nisbett, R. E. (1986). *The effects of statistical training on thinking about everyday problems. Cognitive Psychology*, 18, 253–292.

Geis, M. C., & Zwicky, A. M. (1971). On invited inferences. *Linguistic Inquiry*, 2, 561–566.

Golding, E. (1981). *The effect of past experience on problem solving*. Paper presented at the Annual Conference of the British Psychological Society. Surrey University.

Goodman, N. (1965). *Fact, fiction, and forecast*. Indianapolis: Bobbs–Merrill.

Griggs, R. A. (1983). The role of problem content in the selection task and in the THOG problem. In J. St. B. T. Evans (Ed.), *Thinking and reasoning: Psychological approaches*. London: Routledge & Kegan Paul.

Griggs, R. A. (1984). Memory cueing and instructional effects on Wason's selection task. *Current Psychological Research and Reviews*, 3, 3–10.

Griggs, R. A., & Cox, J. R. (1982). The elusive thematic-materials effect in Wason's selection task. *British Journal of Psychology*, 73, 407–420.

Hempel, C. G. (1965). *Aspects of scientific explanation, and other essays in the philosophy of science*. New York: Free Press.

Henle, M. (1962). On the relation between logic and thinking. *Psychological Review*, 69, 366–378.

Holland, J., Holyoak, K. J., Nisbett, R. E., & Thagard, P. (1986). *Induction: Processes of inference, learning, and discovery*. Cambridge, MA: Bradford Books/MIT Press.

Jepson, C., Krantz, D. H., & Nisbett, R. E. (1983). Inductive reasoning: Competence or skill? *Behavioral and Brain Sciences*, 6, 494–501.

Johnson-Laird, P. N. (1983). *Mental models*. Cambridge, MA: Harvard Univ. Press.

Johnson-Laird, P. N., Legrenzi, P., & Legrenzi, M. S. (1972). Reasoning and a sense of reality. *British Journal of Psychology*, **63**, 395–400.

Kelley, H. H. (1972). Causal schemata and the attribution process. In E. E. Jones et al. (Eds.), *Attribution: Perceiving the causes of behavior*. Morristown, NJ: General Learning Press.

Kelley, H. H. (1973). The process of causal attribution. *American Psychologist*, **78**, 107–128.

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.

Nelson, E. J. (1933). On three logical principles in intension. *Monist*, **43**, 268–289.

Nisbett, R. E., Krantz, D. H., Jepson, C., & Kunda, Z. (1983). The use of statistical heuristics in everyday inductive reasoning. *Psychological Review*, **90**, 339–363.

Piaget, J., & Inhelder, B. (1975). *The origin of the idea of chance in children*. New York: Norton. (Original work published in 1951)

Quine, W. V. O. (1974). *Methods of logic* (3rd ed.). London: Routledge & Kegan Paul.

Reich, S. S., & Ruth, P. (1982). Wason's selection task: Verification, falsification and matching. *British Journal of Psychology*, **73**, 395–405.

Rips, L. J. (1983). Cognitive processes in propositional reasoning. *Psychological Review*, **90**, 38–71.

Rips, L. J., & Conrad, F. G. (1983). Individual differences in deduction. *Cognition and Brain Theory*, **6**, 259–285.

Simco, N. D., & James, G. G. (1982). *Elementary logic* (2nd ed.). Belmont, CA: Wadsworth.

Wason, P. C. (1966). Reasoning. In B. M. Foss (Ed.), *New horizons in psychology I*. Harmondsworth: Penguin.

Wason, P. C., & Johnson-Laird, P. N. (1972). *Psychology of reasoning*. Cambridge, MA: Harvard Univ. Press.

Wason, P. C., & Shapiro, D. (1971). Natural and contrived experience in a reasoning problem. *Quarterly Journal of Experimental Psychology*, **23**, 63–71.

Yachanin, S. A. (1985). *Facilitation in Wason's selection task: Content and instruction*. Unpublished manuscript, Lake Erie College.