

Illusions in quantified reasoning: How to make the impossible seem possible, and vice versa

YINGRUI YANG and P. N. JOHNSON-LAIRD
Princeton University, Princeton, New Jersey

The mental model theory postulates that reasoners build models of the situations described in premises, and that these models normally represent only what is true. The theory has an unexpected consequence. It predicts the existence of *illusions* in inferences. Certain inferences should have compelling but erroneous conclusions. Two experiments corroborated the occurrence of such illusions in inferences about what is possible from disjunctions of quantified assertions, such as, "at least some of the plastic beads are not red." Experiment 1 showed that participants erroneously inferred that impossible situations were possible, and that possible situations were impossible, but that they performed well with control problems based on the same premises. Experiment 2 corroborated these findings in inferences from assertions based on dyadic relations, such as, "all the boys played with the girls."

Logical reasoning is a fundamental human ability, but what underlies it is a matter of controversy. Theorists argue that it depends on rules that have a specific content (Cheng & Holyoak, 1985), on memories for previous cases of reasoning (Kolodner, 1993), on innate modules attuned to specific contents (Cosmides, 1989; Cummins, 1996), and on connectionist networks of associations (Shastri & Ajanagadde, 1990). Each of these possibilities may play a part, but none can explain the whole story, because human reasoners can make deductions about matters of which they have no prior knowledge. For example, if one is told

All the beads are plastic.

All the plastic things are red.

one can infer

All the beads are red.

even if one knows nothing else about the beads. This sort of reasoning cannot be explained by any of the theories cited above. It depends on general principles that concern quantifiers, such as *all*, *some*, and *none*, and sentential connectives, such as *if*, *or*, and *and*. A major part of the current controversy concerns the nature of these general principles.

On the one hand, certain theorists have proposed that the mind is equipped with *formal rules of inference* akin to

those of a logical calculus (see, e.g., Braine & O'Brien, 1991; Rips, 1994). One such rule, for example, has the following form:

All A are B.

All B are C.

∴ All A are C.

This rule would enable one to draw the deduction above, and the system of rules proposed by Rips would allow many deductions to be made, although no current psychological theory based on formal rules is "complete" in the sense that none allows for all possible valid deductions to be drawn.

On the other hand, some theorists have argued that the untrained mind is not equipped with formal rules of inference, but rather relies on the general principle that a valid deduction is one in which the conclusion must be true given that the premises are true (Johnson-Laird, 1983; Johnson-Laird & Byrne, 1991; Polk & Newell, 1995). This principle is put into practice by constructing *mental models* of the premises, formulating a conclusion on the basis of such models—if none is provided by a helpful interlocutor—and checking its validity by examining whether or not it holds in the models of the premises.

A fundamental principle of this model theory, the "principle of truth," is that reasoners *normally* represent only what is true in order to minimize the load on working memory. This principle is subtle because it is applicable at two levels. First, mental models represent only true possibilities. Second, within those true possibilities, they represent the literal propositions (affirmative or negative) in a premise only when they are true. If a proposition is false, its negation is true. But, what reasoners fail to represent is falsity—that is, false affirmative propositions in the premises and false negative propositions in the premises. We can best illustrate the way the principle works

This research was supported in part by ARPA (CAETI) Contracts N66001-94-C-6045 and N66001-95-C-8605 and by the Xian-Lin Ji Foundation of Peking University, China. We thank the members of our laboratory, who provided much helpful advice: Patricia Barres, Victoria Bell, Zachary Estes, Yevgeniya Goldvarg, and Mary Newsome. We are grateful to Monica Bucciarelli, who investigated the construction of external models in quantified reasoning. We also thank Michael Oakford, David O'Brien, Tom Ward, and three anonymous reviewers for their constructive criticisms. Correspondence concerning this article should be addressed to Y. Yang or P. N. Johnson-Laird, Department of Psychology, Green Hall, Princeton University, Princeton, NJ 08544 (e-mail: yingruiy@phoenix.princeton.edu or phil@princeton.edu).

with an example. Consider an exclusive disjunction about a hand of cards that contains two literals, one negative and one affirmative:

There is *not* a king in the hand, or else there is an ace in the hand

At the first level of the principle of truth, reasoners construct two alternative models corresponding to the two true possibilities:

\neg king
ace

where each model is shown on a separate line, and “ \neg ” denotes negation. Each model corresponds to a true possibility, given the disjunction—that is, each model corresponds to a true row in a truth table. At the second level of the principle, each model represents only those *literal* propositions in the disjunctive premise that are true within the possibility. Hence, the first model represents the truth of the negative disjunct in the premise (there is not a king in the hand), but it does not represent the falsity of the affirmative disjunct (there is an ace in the hand). Similarly, the second model represents the truth of the affirmative disjunct in the premise, but it does not represent the falsity of the negative disjunct. Thus, mental models normally represent the literals in the premises when they are true in the true possibilities, but not when they are false.

Reasoners make mental “footnotes” to keep track of what is false. The footnotes make it possible to flesh out the models explicitly. And only *fully explicit* models of what is possible represent both the true and the false literals in each model. Thus, the fully explicit models of the disjunction above are:

\neg king \neg ace
king ace

Fully explicit models can be used to reason in a completely valid way akin to the use of truth tables.

The principle of truth concerns the normal interpretation of assertions, particularly assertions containing logical terms, such as sentential connectives and quantifiers. Clearly, individuals can keep track of what would be false in the case of simple categorical assertions, such as:

It is raining.

At what point reasoners cease to bear falsity in mind is an empirical question, but previous studies suggest that this point is normally reached when a set of premises calls for more than one model—that is, when the premises are compatible with several possibilities (see, e.g., Johnson-Laird & Savary, 1996). In such a case, mental footnotes about falsity are likely to be soon forgotten, and so the predictions presented here depend only on mental models, not on fully explicit models.

One of the authors (J.-L.), while developing a computer implementation of the mental model theory, discovered an odd output to a particular inference. He first thought

there was a bug in the program. When he worked out the correct answer by hand, however, he discovered that the program was correct and that he had succumbed to an illusion. Indeed, the principle of truth has an unexpected consequence. It predicts the existence of inferences that yield illusory conclusions—that is, conclusions that nearly everyone draws, that seem compelling, but that are invalid. After the discovery of potential illusions, Rips (1994) tried his PSYCOP program on one of them, and the program drew no conclusion whatsoever (i.e., neither the correct conclusion nor the illusory conclusion the model theory had predicted). Hence, Rips’s theory, like Braine’s theory (see, e.g., Braine & O’Brien, 1991) does not predict illusions. Yet, illusions do occur, and previous studies have demonstrated this in inferences that are based on sentential connectives (e.g., Johnson-Laird & Savary, 1996). Here is an example of a simple but compelling illusion (see Johnson-Laird & Goldvarg, 1997):

Only one of the following premises is true about a particular hand of cards:

There is a king in the hand or there is an ace, or both.

There is a queen in the hand or there is an ace, or both.

There is a jack in the hand or there is a ten, or both.

Is it possible that there is an ace in the hand?

Nearly everyone responds “yes,” but the response is an illusion because the presence of an ace in the hand renders two of the premises true contrary to the rubric that only one of them is true.

How does the model theory predict that reasoners will err in this way? According to the principle of truth, reasoners fail to represent what is false. The problem above is equivalent to an exclusive disjunction of the form: P or else Q or else R, where only one of the three propositions, each of which is a disjunction, can be true. Hence, when naive reasoners think about the truth of the first disjunction in the problem, they fail to bring to mind the concomitant falsity of the other two disjunctions. To think about the truth of the first disjunction is to consider the possibilities in which it would be true—that is, to call to mind the following three models:

king
ace
king ace

These models imply that an ace is possible, because it occurs in two of them. But, the response is a fallacy. When the first disjunction is true, the second disjunction is false, which establishes that there is neither a queen nor an ace:

\neg queen \neg ace

Reasoners who go on to consider the truth of the second premise are likely to make the same mistake. They will consider its models:

queen

ace
queen ace

and overlook that when it is true, the first disjunction is false and so there is neither a king nor an ace:

¬ king ¬ ace

In summary, naive reasoners will consider the models of the first disjunction and infer that an ace is possible, and if they go on to consider the models of the second disjunction, they will draw the same conclusion. The error arises because they fail to consider the falsity of the other disjunctions when thinking about the truth of one disjunction.

A more powerful sort of reasoning hinges on quantifiers, so in the present studies, we examined illusory inferences based on quantified assertions. In this domain, the *same* premises can be used to elicit both illusory inferences and control inferences for which the theory predicts that the conclusions should be correct even if reasoners fail to consider falsity. In this way, we can eliminate the hypothesis that illusory premises are somehow too difficult for logically untrained individuals to understand. In what follows, we consider the mental models of quantified assertions and then the predictions of the illusions.

How are quantified assertions, such as “at least some of the A are not B,” represented in mental models? The theory distinguishes between the initial models of the premises and those that can be constructed from them through simple manipulations. Bucciarelli and Johnson-Laird (1999) examined the external models that individuals constructed with cut out shapes when they reasoned syllogistically, and we will follow their results, which differed, in part, from the accounts in Johnson-Laird and Byrne (1991, 1996). Reasoners made a variety of different interpretations of quantified statements, but the prin-

cipal data for our purposes are the initial models that they constructed for different sorts of quantified assertions.

Table 1 presents the initial models and the alternative models of each of the four main sorts of singly quantified assertion (see Bucciarelli & Johnson-Laird, 1999). Thus, given a premise of the form *all the A are B*, most of the models represented the two sets as coextensive but a few models represented A as properly included within B. We emphasize that the use of the definite article in the premise is intended to convey that there is no doubt about the existence of As. Our instructions in the experiments also established this point. A model of the coextensive interpretation depicts a small but arbitrary number of individuals who are As and thus Bs, and reasoners needed to make a mental “footnote” that the model represents the set of As in its entirety:

[a] b
[a] b

...

Each row in this diagram represents a separate individual in a single model; the square brackets represent that the set of As has been represented in its entirety, and the ellipsis allows for the existence of other sorts of individuals. It follows that the individuals corresponding to the ellipsis can be Bs, but not As. Reasoners can accordingly construct a model corresponding to the other interpretation of the premise:

[a] b
[a] b
b

in which A is properly included in B.

Table 1
The Initial Mental Models and the Alternative Mental Models of the Four Main Singly-Quantified Assertions

Assertions	Initial Mental Model	Alternative Mental Models
All the A are B:	[a] b [a] b ...	[a] b [a] b b ...
At least some of the A are B:	a b a b ...	a b a b a b a b a b a b a b a b a b a b a b a b a b a b a b a b ...
None of the A is a B:	[a] [a] [b] [b] ...	None None None None ...
At least some of the A are not B:	a a [b] [b] ...	a a a a a a a a a [b] a [b] a [b] a [b] ...

Note—The brackets indicate that a set of individuals has been represented in its entirety (i.e., no more individuals of this sort can be added to the model). The initial model of *some A are not B* supports the converse conclusion, but this conclusion can be refuted by an alternative model of the premise.

Given a premise of the form, *some of the A are B*, the majority of models represented B as properly included within A, but other models represented the two sets as coextensive, or overlapping, or with A properly included in B. Thus, the initial model is of the form:

```

a      b
a
      b
...

```

Because neither set is represented in its entirety, reasoners can add new tokens to the model to construct the alternative interpretations. Given a premise of the form, *none of the A are B*, all the models represented the two sets as disjoint. Given a premise of the form, *some of A are not B*, the majority of models represented B as properly included within A, or as overlapping with A—most participants evidently inferred that some of the A are B (see Grice, 1975), but other models represented the two sets as disjoint. To enhance the possibility of this interpretation in our experiments, we used the form “at least some A are not B,” and the initial model in Table 1 allows for the disjoint interpretation.

Although the construction of external models may not reflect the internal representation of premises, the interpretations in Table 1 are comparable to the findings of other studies (see, e.g., Begg, 1987; Johnson-Laird, 1970), and, given the existence of As and Bs, they correspond to the standard account of the set-theoretic relations of the assertions as represented by Euler circles. Indeed, the models might take the form of Euler circles (see Stenning & Yule, 1997), but the participants in Bucciarelli and Johnson-Laird’s (1999) experiments used individual tokens, which also generalize to capture inferences that depend on relations. Relational inferences cannot be represented by Euler circles. For example, all horses are animals; therefore, all horses’ heads are animals’ heads. Hence, a reasonable assumption is that logically untrained individuals are likely to be biased toward the initial models in Table 1 and yet be capable of constructing the alternative models (see also Polk & Newell, 1995, for the case for alternative interpretations). The alternative models are likely to be important in reasoning about possibilities, as in the present studies. Each model can be constructed merely by adding tokens that represent As or Bs the initial model of an assertion.

How can quantified assertions give rise to illusory inferences? One method depends on combining assertions in exclusive disjunctions. Not all illusory inferences depend on such disjunctions (see, e.g., Johnson-Laird & Savary, 1996), but the illusions in the present experiments are based on them. The principle of truth yields models of an exclusive disjunction of two propositions that consist merely in the models of the first proposition and the models of the second proposition. We exploited this principle to construct *illusions of possibility*—with a ques-

tion about a possibility to which reasoners should tend to respond “yes” when, in fact, the correct answer is “no.” We refer to these problems as “yes/no” problems, an abbreviation that states the predicted answer followed by the correct answer. Here is an example:

(1) Only one of the following statements is true:

At least some of the plastic beads are not red, or

None of the plastic beads are red.

Is it possible that none of the red beads are plastic?

We used the rubric “only one of the following statements is true,” and the connective “or,” in an effort to make clear that one premise was true and the other premise was false. The instructions spelled out this information explicitly. Given the premises of Problem 1, reasoners should consider the two premises separately. They should construct a model of the first premise (see Table 1),

```

p
p
      [r]
      [r]
...

```

where “p” denotes a plastic bead and “r” denotes red. The model supports the possibility that none of the red beads are plastic, and so reasoners should respond “yes.” They may consider a model of the second premise:

```

[p]
[p]
      [r]
      [r]
...

```

This model also supports the response “yes.” Some individuals may not consider the second premise of this problem if they have already responded “yes” on the basis of the first premise. In either case, however, the response is an illusion. When the first premise is true, the second premise is false; that is, some of the plastic beads are red, and so a fully explicit model of this case, representing both truth and falsity, is

```

p      r
p      ¬r
...

```

Conversely, when the second premise is true, the first premise is false; that is, all of the plastic beads are red, which conflicts with the second premise, and therefore, there is no resulting model—strictly speaking, the result

is the null model. The remaining model above refutes the possibility that none of the red beads are plastic, therefore, the correct answer to the question is “no.”

Control problems are those that should elicit the correct response even if reasoners fail to consider falsity. Hence, the controls for illusions of possibility should elicit the correct answer, “yes.” We refer to these problems as “yes/yes” problems. An example, based on the same premises as Problem 1 above, poses a different question:

- (2) Is it possible that at least some of the red beads are plastic?

This should elicit the answer “yes,” because an alternative model of the first premise (see Table 1) is one in which the sets of ps and rs overlap:

```
p
p      [r]
      [r]
...

```

The answer “yes” to this problem is correct, as is shown above by the fully explicit model of the premises. Thus, the failure to consider falsity does not lead to errors with the control problem.

We also examined *illusions of impossibility* in which the mental models support the answer “no,” but the correct answer is “yes.” We refer to these problems as “no/yes” problems. With the preceding premises, the following question

- (3) Is it possible that all of the red beads are plastic?

is such a case. Neither of the two premises has a mental model that supports this conclusion, and so the participants should respond “no.” But, as the fully explicit model above shows, when falsity is taken into account, the correct answer is “yes.” And again, where possible, we used the same premises to generate a control problem for illusions of impossibility in which the correct answer was “no.” We refer to these problems as “no/no” problems. With the preceding premises, the following question is such a case:

- (4) Is it possible that all the plastic beads are red?

Neither of the two premises has a model that supports this conclusion, but in this case the correct answer is “no” (as shown by the fully explicit model). Our first experiment tested whether logically untrained individuals succumb to illusions in quantified inferences about possibilities, but make correct responses to the control problems.

EXPERIMENT 1

Method

Design. The participants acted as their own controls and carried out four sorts of modal inference concerning what was possible: illusions of possibility (yes/no problems), controls for illusions of possibility (yes/yes problems), illusions of impossibility (no/yes

problems), and controls for illusions of impossibility (no/no problems). The inferences were made on the basis of five pairs of monadic premises that each referred to the same two terms. The five pairs of premises were combined on separate trials with four different modal conclusions, which made a total of 20 different problems. Each of these problems had the following form:

Only one of the following statements is true:

Premise 1, or

Premise 2

Is it possible that ...?

For example, given the premises:

Only one of the following statements is true:

At least some of the plastic beads are not red, or

None of the plastic beads are red.

the four sorts of problems were created by presenting the following questions on separate trials:

- (1) It is possible that none of the plastic beads are red? (illusion of possibility: yes/no.)
- (2) Is it possible that at least some of the red beads are plastic? (control: yes/yes.)
- (3) Is it possible that all the red beads are plastic? (illusion of impossibility: no/yes.)
- (4) Is it possible that all of the plastic beads are red? (control: no/no.)

In the experiment, however, each of these problems was presented with a different lexical content.

In order to develop the set of problems, we examined all possible pairs of quantified statements in which both premises referred to the same terms; for example:

All the A are B.

Some of the A are B.

We excluded pairs that had the same statement twice, pairs that differed merely in the order of the two premises, and the variants created by the symmetric premises (*some of the A are B/some of the B are A*, and *none of the A are B/none of the B are A*). Five of the remaining pairs yielded viable problems. Two of the pairs yielded all four types of problem when they were combined with different conclusions (as in the example above). The three remaining pairs yielded at least two or more of the four sorts of problem. Ten of the problems were illusions, and 10 of them were controls. Table 2 presents the full set of problems, their mental models, their fully explicit models, and the four questions for each pair with their predicted and correct answers. The 20 problems had a different lexical content, and they were presented in one of six different random orders with approximately an equal number of participants receiving each order.

Materials. The content of the problems concerned beads of different colors (blue, red, green, or brown), shapes (square, round, triangular, or rectangular), and substances (wooden, plastic, metal, or cement). We chose among these properties at random to make 20 sets of materials, and then assigned them at random to the problems.

Procedure. The participants were tested individually in a quiet room. They were given two booklets. The first booklet provided the instructions and a practice problem. The experimenter read the instructions aloud while the participants followed in the booklet. The instructions stated that the participants' task was to answer a series of questions about various possibilities on the basis of information that they would receive about beads. The instructions made clear that all the different sorts of beads existed within the hypothetical domain of the experiment:

Table 2
The Premises, Their Mental Models and Fully Explicit Models,
and the Four Questions and Their Predicted and Correct Answers for Experiment 1

Premises and Questions	Mental Models	Fully Explicit Models
Only one is true:	a	[a] a -b
Some A are not B.	a	[a] a b
No A are B.	[b]	[b]
	[b]	[b]
1. Possible that no B are A?		Illusion of possibility: Yes/No
2. Possible that some B are A?		Control for "Yes": Yes/Yes
3. Possible that all B are A?		Illusion of impossibility: No/Yes
4. Possible that all A are B?		Control for "No": No/No
Only one is true:	a	b [a] b a b
Some A are B.	a	[a] b a -b
All A are B.	b	
5. Possible that all A are B?		Illusion of possibility: Yes/No
6. Possible that some B are A?		Control for "Yes": Yes/Yes
7. Possible that all B are A?		Illusion of impossibility: No/Yes
8. Possible that no A are B?		Control for "No": No/No
Only one is true:	b	a b [b] a [a] -b
Some B are not A.	b	a [b] a [a] -b
Some A are B.	[a]	b [a] b [b]
	[a]	[a] [b]
9. Possible that all B are A?		Control for "Yes": Yes/Yes
10. Possible that some A are not B?		Control for "Yes": Yes/Yes
11. Possible that all A are B?		Illusion of impossibility: No/Yes
12. Possible that no A are B?		Illusion of impossibility: No/Yes
Only one is true:	a	b [b] a [a] b
Some A are not B.	a	b [b] a [a] b
Some B are not A.	[b]	[a] a b
	[b]	[a] a b
13. Possible that no A are B?		Illusion of possibility: Yes/No
14. Possible that some A are B?		Control for "Yes": Yes/Yes
15. Possible that all A are B?		Illusion of impossibility: No/Yes
16. Possible that all B are A?		Illusion of impossibility: No/Yes
Only one is true:	[a]	b [b] a [a] b a
All A are B.	[a]	b [b] a [a] b a
All B are A.		b a b
17. Possible that all A are B?		Control for "Yes": Yes/Yes
18. Possible that some A are not B?		Illusion of impossibility: No/Yes
19. Possible that no A are B?		Control for "No": No/No
20. Possible that no B are A?		Control for "No": No/No

Note—"Yes"/"No" indicates that the predicted answer is "Yes," but the correct answer is "No." We have omitted the implicit models.

All the problems concern the following situation. First, assume that there is a group of children and each of them has a bag. . . . Second, imagine that there are many beads and that the manufacturer puts them in bags. The beads may vary in color . . . , shape . . . , and material

The instructions also made clear how to interpret the initial statement in each problem that only one of the two premises was true: "You will notice that every problem contains two statements, but only one of them is true, i.e., one is true and the other is false, though you do not know which of them is true." The participants were also told to write down a "yes" or a "no" to each problem, and then to rate their confidence in their answer, using a 5-point scale, in which 1 indicated *no confidence* and 5 indicated *complete confidence*." They were told that they could take as much time as they wanted to complete the task. After they had asked any questions about the task, they were given a simple yes/yes problem for practice:

- Only one of the following statements is true:
- At least some of the brown beads are round, or
- All the brown beads are round.
- Is it possible that at least some of the brown beads are round?

Once the participants understood the task, particularly the statement that one premise was true and one premise was false, the experiment was begun. The problems were presented in a second booklet, with each problem on a separate page.

Participants. Twenty undergraduate students took part in the experiment in order to fulfill a requirement of an introductory psychology course at Princeton University. None had received any previous training in formal logic or had participated in an experiment on reasoning.

Table 3
The Percentages of Correct Responses to the Four Sorts of Problems in Experiment 1

	Illusions	Controls
Inferences of possibility	23 (4.35)	87 (4.43)
Inferences of impossibility	67 (4.36)	94 (4.71)
Overall	54 (4.39)	90 (4.51)

Note—The figures in parentheses are the participants' mean confidences in their answers (on a five-point scale where 1 signifies *no confidence* and 5 signifies *full confidence*).

Table 4
The Percentages of Correct Responses to Each of the Twenty Problems in Experiment 1
and to Each of the Related Problems in Experiment 2,
in Which the Premises Were Dyadic in Form

Premises	Status of Question	Percentages of Correct Responses	
		Experiment 1	Experiment 2
Only one is true:			
Some A are not B.			
No A are B.			
1. Possible that no B are A?	Yes/No	20	5
2. Possible that some B are A?	Yes/Yes	80	95
3. Possible that all B are A?	No/Yes	70	67
4. Possible that all A are B?	No/No	95	95
Only one is true:			
Some A are B.			
All A are B.			
5. Possible that all A are B?	Yes/No	25	14
6. Possible that some B are A?	Yes/Yes	100	100
7. Possible that all B are A?	No/Yes	70	90
8. Possible that no A are B?	No/No	100	95
Only one is true:			
Some B are not A.			
Some A are B.			
9. Possible that all B are A?	Yes/Yes	60	81
10. Possible that some A are not B?	Yes/Yes	100	100
11. Possible that all A are B?	No/Yes	80	71
12. Possible that no A are B?	No/Yes	55	67
Only one is true:			
Some A are not B.			
Some B are not A.			
13. Possible that no A are B?	Yes/No	25	19
14. Possible that some A are B?	Yes/Yes	95	95
15. Possible that all A are B?	No/Yes	55	62
16. Possible that all B are A?	No/Yes	60	33
Only one is true:			
All A are B.			
All B are A.			
17. Possible that all A are B?	Yes/Yes	90	90
18. Possible that some A are not B?	No/Yes	80	76
19. Possible that no A are B?	No/No	95	95
20. Possible that no B are A?	No/No	85	86

Results and Discussion

The overall percentages of correct responses (and the confidence ratings) for each of the four sorts of inference are presented in Table 3. The percentages of correct responses for each of the 20 problems are presented in Table 4. The participants were correct on 90% of the control inferences, but were correct only slightly more than chance (54%) on the illusory inferences. All the participants were more accurate with the control inferences than with the illusory ones ($p = .52^0$, which is less than 1 in a million). They were also slightly more confident in their answers to the control problems than in their answers to the illusory problems. This difference—a mean of .13 on the 5-point scale—was reliable (Wilcoxon's test, $N = 12$, $T+ = 65$, $p < .03$). The inferences of impossibility yielded a greater percentage of correct responses than the inferences of possibility (Wilcoxon's test, $z = 2.72$, $p < .01$). In addition, there was a reliable interaction. The difference between the control problems and the illusions was greater for the inferences of possibility than for the inferences of impossibility (Wilcoxon's test, $z = 2.77$, $p < .01$).

This interaction shows that illusions of possibility are much more convincing than illusions of impossibility—an effect that we have corroborated in modal inferences that depend on sentential connectives rather than on quantifiers (see Johnson-Laird & Goldvarg, 1997). We propose an explanation of this difference later.

The overall pattern of results corroborates the predictions made on the basis of the mental model theory. Reasoners were more accurate in the control inferences than in the illusory inferences. But, as Table 4 shows, there was also some variation in difficulty, particularly within the illusions of impossibility and the control problems of possibility. Before we could give an account of this variation, we needed to determine whether or not it was robust, and whether the general pattern of results would be replicable; therefore, we conducted a second experiment.

EXPERIMENT 2

Experiment 1 confirmed the existence of illusions that are based on quantified statements in the case of simple

monadic premises, such as “all the plastic beads are red.” We needed to corroborate the existence of illusions and to extend them to assertions that were syntactically and semantically different from those in Experiment 1. Hence, we examined a new set of inferences in which the premises concerned dyadic relations—for example, “all the boys got red beads.” In these problems, the premises contained a main verb that described a relation, such as “got,” that holds between two sets of individuals.

Method

Design. The participants acted as their own controls and carried out four sorts of modal inference that concerned what was possible. These were based on five pairs of dyadic premises that each referred to the same two terms. These five pairs were combined on separate trials with four different modal conclusions, making a total of 20 different inferences. These were equivalent to the problems used in Experiment 1, except that the premises were dyadic. The 20 problems had a different lexical content and were presented in one of six different random orders with approximately an equal number of participants receiving each order.

Materials. Each pair of dyadic premises referred to two sets of individuals or entities, such as boys, girls, or books; they concerned the five relations between one set and the other—that is, “played with,” “liked,” “found,” “got,” and “collected.”

Procedure. The participants were tested individually. The procedure and instructions were identical to those in Experiment 1 except that the monadic premises were rephrased by the dyadic ones.

Participants. Twenty undergraduate students from the same population as in Experiment 1 participated in the experiment.

Results

The percentages of correct responses to each of the 20 problems are presented in Table 4. The overall percentages of correct responses (and the confidence ratings) for each of the four sorts of inference are summarized in Table 5. Again, all of the participants were more accurate with the control inferences (93% correct) than with the illusory ones (51% correct; $p = .52^0$, which is less than 1 in a million). They were slightly more confident in their answers to the control problems than in their answers to the illusory problems. The mean difference of .32 on the 5-point scale was reliable (Wilcoxon’s test, $z = 3.62, p < .001$). The inferences of impossibility yielded a greater percentage of correct responses than did the inferences of possibility (Wilcoxon’s test, $z = 1.82, p < .05$). As before, there was a reliable interaction. The difference between the control problems and the illusions was greater for the inferences of possibility than for the inferences of impossibility (Wilcoxon’s test, $z = 3.22, p < .001$). Illusions of possibility are much more convincing

than illusions of impossibility, as we observed in Experiment 1. This difference relates to another observation that we made in Experiment 1—that is, there was some variation in difficulty, particularly within the illusions of impossibility and the control inferences of possibility. Experiment 2 confirmed this variation. A scatter plot of the percentages of correct responses between each form of problem in Experiments 1 and 2 is presented in Figure 1. The correlation between the results of the two experiments was very high (Kendall’s tau = .94, $p < .0001$). Evidently, the pattern of results with these quantified inferences is highly robust.

GENERAL DISCUSSION

Our experiments showed that individuals succumb to illusions in inferences about what is possible, and what is impossible, when they are given quantified assertions about individuals. The phenomena were predicted by the model theory’s principle of truth according to which reasoners take into account what is true, but not what is false. For example, given the following problem,

Only one of the following statements is true:

At least some of the metal beads are not blue, or

At least some of the blue beads are not metal.

Is it possible that none of metal beads are blue?

three quarters of the participants responded, “yes,” even though the correct response is “no.” We emphasize that the instructions to the participants made clear that the rubric to the problem meant that one of the premises was true and the other was false, although there was no way to know which was which. The model theory predicts that reasoners infer that the questioned conclusion is possible given the first premise and that they fail to take into account the falsity of the second premise in such a case. Accordingly, participants construct the following sort of model to represent that the first premise is true:

```
metal
metal
    [blue]
    [blue]
    ...
```

This model is consistent with the putative conclusion, and so they respond “yes.” If they had taken into account the concomitant falsity of the second premise, they would have constructed the following model:

```
metal
metal
metal [blue]
metal [blue]
    ...
```

Table 5
The Percentages of Correct Responses to the Four Sorts of Problems in Experiment 2

	Illusions	Controls
Inferences of possibility	13 (4.37)	94 (4.31)
Inferences of impossibility	70 (4.12)	93 (4.69)
Overall	51 (4.22)	93 (4.54)

Note—The figures in parentheses are the participants’ mean confidences in their answers (on a five-point scale, where 1 signifies *no confidence* and 5 signifies *full confidence*).




Figure 1. A scatter plot of the percentages of correct responses for each form of problem in Experiments 1 and 2.

This satisfies both the truth of the first premise and the falsity of the second premise, and they would have responded correctly “no.” The participants who consider instead the truth of the second premise would fall into the same trap by failing to consider the falsity of the first premise.

Is there an alternative explanation of our results? One hypothesis is that the instructions, materials, or task, were too difficult for the participants to understand. Two phenomena, however, argue against this idea. First, the participants performed very well with the control problems (over 90% correct in each experiment). These control problems were based on logically identical premises to those used for the illusions; they differed only in the question posed to the participants. These questions themselves varied from one problem to another, so that, for example, a question of the form

Is it possible that none of the A are B?

occurred sometimes for an illusion and sometimes for a control problem (see Table 2). Hence, there was unlikely to be anything about the instructions, materials, or task that was intrinsically difficult to understand. Second, the participants were highly confident in their responses, albeit slightly more confident in their responses to the con-

trols than to the illusions, with mean ratings for the illusions of 4.39 (Experiment 1) and 4.22 (Experiment 2) on a 5-point scale. These ratings suggest that the participants had no major difficulties with the task or the materials. They may have sensed, however, that something was not quite right about their illusory responses, because they were slightly less confident in them.

Another hypothesis is that the participants made their inferences on the basis of whether or not the quantified assertion in the conclusion was identical to one of the premises. For example, they responded “yes” to Problem 5 (in Table 2) because the question was of the form “is it possible that all the A are B;” and the second premise was of the form “all the A are B.” Although we cannot rule out the possibility that such identities might sometimes have influenced performance, they cannot account for the differences between the illusions and the control problems. If the participants had responded “yes” whenever there was such an identity, and “no” otherwise, the pattern of their responses would have been wholly different from those that we observed. For example, the participants would have responded “no” uniformly to Problems 1–4.

Is it possible that the participants misinterpreted the quantified assertions, which somehow justified the illu-

sory responses? Indeed, is it possible that the illusions are simply a by-product of well known reasoning errors, such as the invalid conversion of premises? Once again, such a hypothesis can explain some illusions, but cannot explain all of them. Consider, for example, the following case, which concerns Problem 1:

Only one of the following statements is true:

At least some A are not B, or

None of the A are B.

Is it possible that none of the B are A?

If the second premise is true, so is the first. But, in this case, the rubric that only one of the statements is true, is wrong. Reasoners can, of course, consider the truth of the first premise, which does not imply that the second premise is true. But, it is possible that reasoners overlook this case, and instead search for an interpretation of the first premise that is not entailed by the truth of the second premise. They might accordingly reinterpret the first premise so that it means:

Some but not all A are not B.

The truth of the second premise is now compatible with the falsity of the first premise. Hence, it follows that the correct answer to the question is “yes.” This hypothesis, however, fails with other problems. Consider Problem 13:

Only one of the following statements is true:

At least some A are not B.

At least some B are not A.

Is it possible that no A are B?

The truth of the first premise does not imply the truth of the second premise, and the truth of the second premise does not imply the truth of the first premise. Indeed, individuals readily construct a model of *some A are not B* in which B is properly included within A (Bucciarelli & Johnson-Laird, 1999). This interpretation ought to lead the participants to respond “no.” In fact, they succumb to the illusion of possibility and respond “yes.”

Each of the alternative hypotheses can explain some of the illusions, but not all of them. In contrast, the principle of truth predicts all of the illusions. Moreover, even if the mental model theory was wrong and reasoners represented premises using Euler circles, they would still go wrong if they failed to take falsity into account. The principle of truth also accounts for the finding that illusions of possibility were more compelling than illusions of impossibility and for the variation in the results of the illusions of impossibility and in the control inferences of possibility. The explanation is simple. Illusions of possibility (yes/no) occur, by definition, where the questioned statement is possible when one takes into account the truth of

one of the premises and ignores the falsity of the other. At the other extreme, the control inferences of impossibility (no/no) are those in which the questioned statement is not possible when one takes into account the truth of either of the two premises. But, now consider the other two sorts of inference.

Illusions of impossibility (no/yes) can be divided into two classes. First, there are illusions of possibility in which the questioned statement is possible in at least one model of one of the premises (as shown in Table 1), but not possible in any of the models of the other premise (Problems 3, 12, 15, 16, and 18). Consider Problem 3:

Only one of the following statements is true:

At least some of the A are not B.

None of the A are B.

Is it possible that all the B are A?

The first premise has models in which all the B are A (see Table 1), but the second premise has no such model. Second, there are illusions of impossibility in which the questioned statement is possible in an alternative model of either of the two premises (Problems 7 and 11). Thus, consider Problem 11:

Only one of the following statements is true:

At least some of the B are not A.

At least some of the A are B.

Is it possible that all the A are B?

The first premise has an alternative model (see Table 1, interchanging A and B) in which all the A are B; the second premise does as well. As one would expect, the first sorts of problems are more likely to elicit illusions of impossibility than the second sorts of problems. If we combine the data from the two experiments, there were 62% correct responses for the problems in which the questioned statement is possible in an alternative model of only one premise, and 80% correct responses for the problems in which the questioned statement is possible in an alternative model of either premise (Wilcoxon's test, $z = 4.43$, $p < .001$).

An analogous argument yields two classes of control inferences of possibility (yes/yes). First, there are inferences in which the questioned statement is possible given a model of one of the premises, but not the other (Problems 2 and 9), such as Problem 9:

Only one of the following statements is true:

At least some of the B are not A.

At least some of the A are B.

Is it possible that all the B are A?

Second, there are those in which the questioned statement is possible given an alternative model of either of the premises (Problems 6, 10, 14, and 17), such as Problem 10:

Only one of the following statements is true:

At least some of the B are not A.

At least some of the A are B.

Is it possible that some of the A are not B?

As one would expect, the percentages of correct “yes” responses are higher for the problems in which the conclusion is possible in a model of either premise (96% correct) than for the problems in which the conclusion is possible in a model of only one premise (79% correct; Wilcoxon’s test, $z = 3.01$, $p < .01$). Thus, a more refined analysis based on the model theory can explain some of the variation in performance. We can express these ideas in a slightly different way. The greater the number of models consistent with a conclusion, the more likely the participants were to respond “yes.” If none of the models were consistent with the conclusion, the greater the number of models, then the more likely the participants were to respond “no.” We illustrate the procedure with the four problems based on the following premises:

Only one of the following statements is true:

At least some A are not B.

No A are B.

As Table 1 shows, the first premise is consistent with models in which (i.) the two sets are disjoint, (ii.) the two sets overlap, and (iii.) B is properly included within A. The second premise is consistent with one model (iv.) in which the two sets are disjoint. Given the question

(1) Is it possible that no B are A? (yes/no)

two models (i. and iv.) are consistent with the putative conclusion. Given the question

(2) Is it possible that some B are A? (yes/yes)

two models (ii. and iii.) are consistent with the putative conclusion. Given the question

(3) Is it possible that all B are A? (no/yes)

one model (iii.) is consistent with the putative conclusion. And given the question

(4) Is it possible that all A are B? (no/no)

none of the models is consistent with the putative conclusion. It seems that the participants are following the principle of truth: They consider the set of possible models for one premise without taking into account the falsity of the other premise.

According to the principle of truth, reasoners usually fail to represent what is false. They do so, for example, whether a disjunction is conveyed with a sentential connective, such as

If there is a king in the hand then there is an ace in the hand, or else if there isn’t a king in the hand then there is an ace in the hand.

or with an explicit rubric of the sort used in the present experiments:

Only one of the following assertions is true:

If there is a king in the hand then there is an ace in the hand.

If there is not a king in the hand then there is an ace in the hand.

In each case, participants tend to infer that there is an ace in the hand, which is an illusion. As we mentioned earlier, they can probably cope with a disjunction of simple categorical assertions, such as the following example, which we owe to a reviewer:

Only one of the following assertions is true:

John likes cake.

John likes chocolate flavored food.

Is it possible that John likes chocolate flavored cake?

In this case, they should be able to infer the correct “no” response, because each disjunct calls for merely a single model. They are likely to have difficulties, however, as soon as one of the disjuncts requires them to envisage more than one possibility—that is, to construct more than one model. Here is an example, which is also the simplest illusion that we have so far discovered:

If there is a king in the hand then there is an ace in the hand, or else there is an ace in the hand.

There is a king in the hand.

What follows?

Most people infer wrongly that there is an ace in the hand. One potential antidote to such illusions is to use a rubric that emphasizes falsity. For example, only one of the following assertions is *false*. This yielded a small but reliable improvement in performance in a study by Tabossi, Bell, and Johnson-Laird, (1999).

Readers may wonder what would happen if reasoners were told that *both* premises were false, for example:

Both of the following premises are false:

All the A are B.

Some of the A are B.

In fact, as the principle of truth implies, naive reasoners do not have direct access to the situations that render such assertions false; they have to infer them. They are likely first to infer what it would mean for the first premise to be false and to represent that possibility, and then to infer what it would mean for the second premise to be false and to represent that possibility. They are then likely to try to combine the two cases. The difficulty in running such an experiment, however, is that naive reasoners often err in inferring the situations that render an assertion false; for example, many of them infer that “some of the A are B” is false when “some of the A are *not* B” (see Bucciarelli & Johnson-Laird, 1999).

Other studies of illusions in sentential and probabilistic reasoning have examined other sorts of connectives apart from exclusive disjunctions of premises. For example, Johnson-Laird and Savary (1996) examined the following illusory inference:

If one of the following assertions is true about a particular hand of cards then so is the other:

There is a king in the hand if and only there is an ace in the hand.

There is a king in the hand.

Which is more likely to be in the hand: the king or the ace?

Most of the participants (90%) inferred that the two cards were equally likely to be in the hand. They evidently considered the possibility that both premises were true:

king ace

but overlooked the possibility that they were both false:

¬king ace

Critics have suggested that the model theory is merely postulating that when reasoners think about one disjunct, they forget about the other. They suggest that if the model theory can have such a principle, theories based on formal rules can as well. In fact, the mental model theory has no such principle. It postulates instead that reasoners represent only what is true. This principle implies that when one thinks about one *disjunct*, one ceases to think about the other. But, the principle of truth underlies what is going on. If one merely ceased to think about one proposition and then thinks about another, performance would not be affected by different rubrics, such as the biconditional rubric in the preceding problem. What underlies all the illusions that we have investigated is the failure to consider falsity. The principle of truth is therefore highly plausible. Logically untrained individuals tend to reason on the basis of what is true, not on what is false. In hindsight, some well-known phenomena in reasoning appear to be further manifestations of the principle of truth. An example is the difficulty of inferences in the form known as *modus tollens* (*If A then B, not-B, ∴ not-A*). The illusions are thus interesting and robust shortcomings in human reasoning. They also, however, have theoretical implications.

No current theory of reasoning, apart from the model theory, can account for the illusions. They are beyond the scope of theories that are based on rules that have a specific or pragmatic content (e.g., Cheng & Holyoak, 1985), theories based on a memory for previous cases of reasoning (e.g., Kolodner, 1993), and evolutionary or connectionist theories based on content-specific principles (e.g., Cosmides, 1989; Cummins, 1996; Shastri & Ajjanagadde, 1990). Indeed, it is hardly fair to criticize these theories for failing to explain the illusions, because they are not even intended to account for reasoning about ma-

terials that do not elicit knowledge of a specific domain—that is, for the general ability to reason. In contrast, theories based on formal rules of inference do have this goal (see, e.g., Braine & O'Brien, 1991; Rips, 1994), but they also cannot account for the illusions. These theories apply valid rules of inference to the logical forms of premises, and so they cannot predict the systematic occurrence of invalid conclusions.

Is it possible to save the formal rule theories? One solution is trivial. Formal rules, in the broadest sense, have universal Turing machine power, and so they are equivalent to a programming language (Rips, 1994). Hence, they can be used to simulate any theory whatsoever, including the model theory. It is this principle, of course, that we exploited in implementing the model theory in a computer program. But, in this broad sense, formal rules are no more refutable than a programming language such as LISP. Formal rule theories of reasoning, however, concern rules in a narrower sense. They make use of rules of inference, such as

All A are B.

All B are C.

∴ All A are C.

for drawing inferences by virtue of the logical forms of sentences in natural language. Such rules do not embody the principle of truth.

One possible way to save rule theories may be to argue that the metalinguistic statement “only one of the following statements is true,” somehow leads reasoners to ignore the other statement. But, any argument about the special properties of metalinguistic statements is weak, because illusions work equally well when the metalinguistic statement, “only one of the following statements is true,” is replaced by a simple exclusive disjunction, expressed by “or else,” between the two premises (see Johnson-Laird & Savary, 1996).

Another approach might be an account in which a suppositional strategy is somehow misapplied to yield the illusions. This approach deserves a more detailed comment. Braine's (1998) theory of predicate logic distinguishes between direct reasoning and indirect reasoning. Direct rules (or “schemas”) are supposed to be errorless in practice, whereas indirect inference rules are supposed to be more demanding and to lead to many errors. Direct rules apply to the surface structure of sentences. Indirect rules concern such matters as making suppositions and discharging them. Recent empirical studies of Braine's theory (Yang, 1996; Yang, Braine, & O'Brien, 1998) have investigated only direct reasoning with quantified assertions. The reasoning problems in the present studies are beyond the scope of direct rules. Their inferential structure is relevant to indirect reasoning, and, in particular, to the rule of the “supposition of alternatives.” Consider Problem 1 above, which has the following form:

Only one of the following statements is true:

At least some A are not B, or

No A are B.

Is it possible that no B are A?

According to the rule of supposition of alternatives, reasoners consider the implications of one premise and then the implications of another premise—that is, the rule is appropriate to the *conjunction* of two premises. Both Bonatti and O'Brien (April, 1997 and September 29th, 1997, personal communications) have suggested that reasoners might use the rule inappropriately with disjunctions—that is, they fail to take into account the falsity of other premises. Hence, they make a supposition of the first premise *at least some A are not B*, but cannot infer *no B are A*, because there is no rule that yields this conclusion. But, when they make a supposition of the second premise, *no A are B*, there is a rule that yields *no B are A*. Hence, they respond "yes" to the question. It seems that this simple misapplication of a rule might enable mental logic theory to account for illusory inferences. Unfortunately, the argument fails with other problems, such as Problem 2, which has the same premises as Problem 1, but a different question. The putative conclusion, *at least some B are A*, cannot be inferred from a supposition of either premise. Hence, reasoners should respond "no" to the question, but in fact 80% responded "yes" in Experiment 1 and 95% responded "yes" in Experiment 2. As the mental model theory predicts, Problem 1 is a yes/no illusion and Problem 2 is a yes/yes control. The proposed modification of the "supposition of alternatives," which invalidly applies it to exclusive disjunctions, obliterates the distinction between illusions and controls. It predicts that reasoners will sometimes go wrong with both. In general, modal reasoning with quantifiers is outside the explanatory scope of current mental logic theories.

If we go beyond current theories, other possibilities exist to save the mental logic framework. New formal rules could be introduced to derive conclusions about what is possible or consistent with the premises. A more extreme method is to introduce invalid rules that deliver the illusory conclusions (see Jackendoff, 1988). We are currently exploring these possibilities and at present can report only that it is not easy to preserve formal rules in the face of the illusions. The problem is to ensure that the system of rules yields both illusory inferences and correct conclusions to control problems.

Formal rules and mental models do not exhaust the domain of possible theories of reasoning. Neither sort of theory tells the whole story about reasoning, and some new theory may well transcend them. Our results suggest that this new theory will have to accommodate the principle of truth. Some psychologists have argued that human reasoners never make logical mistakes (see, e.g., Henle, 1962), and others have argued that apparent errors can be explained as a rational performance of a different task from the one the experimenter intended (see,

e.g., Oaksford & Chater, 1994). Our view is that people are rational in principle, but err in practice—as shown, for example, by the ability of the participants to recognize their errors once they are explained to them. Theories of reasoning need to account for logical competence and explain which errors occur and why they occur. The model theory bases rationality on the principle of validity—that is, the ability to grasp that a counterexample to an inference refutes it. This accounts for error in terms of the limitations of working memory, which imply that reasoners are unlikely to cope with fully explicit models. A sensible compromise is for reasoners to take into account what is true, but to treat what is false as though it did not exist. Truth is more useful than falsity, but if reasoners rely only on the truth, they will occasionally succumb to the illusion that they grasp a set of possibilities that is in fact beyond them.

REFERENCES

- BEGG, I. (1987). Some. *Canadian Journal of Psychology*, **41**, 62-73.
- BRAINE, M. D. S. (1998). Steps towards a mental predicate logic. In M. D. S. Braine & D. P. O'Brien (Eds.), *Mental logic* (pp. 273-331). Mahwah, NJ: Erlbaum.
- BRAINE, M. D. S., & O'BRIEN, D. P. (1991). A theory of if: A lexical entry, reasoning program, and pragmatic principles. *Psychological Review*, **98**, 182-203.
- BUCCIARELLI, M., & JOHNSON-LAIRD, P. N. (1999). Strategies in syllogistic reasoning. *Cognitive Science*, **23**, 247-303.
- CHENG, P. W., & HOLYOAK, K. J. (1985). Pragmatic reasoning schemas. *Cognitive Psychology*, **17**, 391-416.
- COSMIDES, L. (1989). The logic of social exchange: Has natural selection shaped how humans reason? Studies with the Wason selection task. *Cognition*, **31**, 187-276.
- CUMMINS, D. D. (1996). Evidence for the innateness of deontic reasoning. *Mind & Language*, **11**, 160-190.
- GRICE, H. P. (1975). Logic and conversation. In P. Cole & J. L. Morgan (Eds.), *Syntax and semantics: Vol. 3. Speech acts* (pp. 41-58). New York: Seminar Press.
- HENLE, M. (1962). On the relation between logic and thinking. *Psychological Review*, **69**, 366-378.
- JACKENDOFF, R. (1988). Exploring the form of information in the dynamic unconscious. In M. J. Horowitz (Ed.), *Psychodynamics and cognition* (pp. 3-10). Chicago: University of Chicago Press.
- JOHNSON-LAIRD, P. N. (1970). The interpretation of quantified sentences. In W. J. M. Levelt & G. B. Flores d'Arcais (Eds.), *Advances in psycholinguistics* (pp. 347-372). Amsterdam: North-Holland.
- JOHNSON-LAIRD, P. N. (1983). *Mental models: Towards a cognitive science of language, inference and consciousness*. Cambridge, MA: Harvard University Press.
- JOHNSON-LAIRD, P. N., & BYRNE, R. M. J. (1991). *Deduction*. Hillsdale, NJ: Erlbaum.
- JOHNSON-LAIRD, P. N., & BYRNE, R. M. J. (1996). Mental models and syllogisms. *Behavioral & Brain Sciences*, **19**, 543-546.
- JOHNSON-LAIRD, P. N., & GOLDBERG, Y. (1997). How to make the impossible seem possible. In *Proceedings of the Nineteenth Annual Conference of the Cognitive Science Society* (pp. 354-357). Mahwah, NJ: Erlbaum.
- JOHNSON-LAIRD, P. N., & SAVARY, F. (1996). Illusory inferences about probabilities. *Acta Psychologica*, **93**, 69-90.
- KOLODNER, J. (1993). *Case-based reasoning*. San Mateo, CA: Morgan Kaufmann.
- OAKSFORD, M., & CHATER, N. (1994). A rational analysis of the selection task as optimal data selection. *Psychological Review*, **101**, 608-631.
- POLK, T. A., & NEWELL, A. (1995). Deduction as verbal reasoning. *Psychological Review*, **102**, 533-566.
- RIPS, L. J. (1994). *The psychology of proof*. Cambridge, MA: MIT Press.

- SHASTRI, L., & AJJANAGADDE, V. G. (1990). An optimal limited inference system. In *Proceedings of the Eighth National Conference on Artificial Intelligence*. Cambridge, MA: MIT Press.
- STENNING, K., & YULE, P. (1997). Image and language in human reasoning: A syllogistic illustration. *Cognitive Psychology*, **34**, 109-59.
- TABOSSI, P., BELL, V., & JOHNSON-LAIRD, P. N. (1999). Mental models in deductive, modal, and probabilistic reasoning. In C. Habel & G. Rickheit (Eds.), *Mental models in discourse processing*. Berlin: John Benjamins.
- YANG, Y. (1996). *Predicting the difficulty of predicate-logic inference problems: An empirical examination of a predicate-logic model*. Unpublished doctoral dissertation, New York University.
- YANG, Y., BRAINE, M. D. S., & O'BRIEN, D. P. (1998). Some empirical justifications of the mental predicate-logic model. In M. D. S. Braine & D. P. O'Brien (Eds.), *Mental logic* (pp. 334-365). Mahwah, NJ: Erlbaum.

(Manuscript received January 27, 1998;
revision accepted for publication January 14, 1999.)