

Illusions in modal reasoning

YEVGENIYA GOLDVARG and P. N. JOHNSON-LAIRD
Princeton University, Princeton, New Jersey

According to the mental model theory, models represent what is true, but not what is false. One unexpected consequence is that certain inferences should have compelling, but invalid, conclusions. Three experiments corroborated the occurrence of such illusions in reasoning about possibilities. When problems had the heading "Only one of the premises is true," the participants considered the truth of each premise in turn, but neglected the fact that when one premise is true, the others are false. When two-premise problems had the heading "One of the premises is true and one is false," the participants still neglected the falsity of one of the premises. As predicted, however, the illusions were reduced when reasoners were told to check their conclusions against the constraint that only one of the premises was true. We discuss alternative explanations for illusory inferences and their implications for current theories of reasoning.

The mental model theory postulates that reasoning is a process in which individuals use the meaning of the premises, together with their general knowledge, to construct mental models of the relevant situations (see, e.g., Johnson-Laird & Byrne, 1991). Each mental model represents a possibility: Its content and structure capture what is common to the different ways in which the possibility may occur. Hence, a conclusion is *possible* if it holds in at least one of the models of the premises; and it is *necessary* if it holds in all the models of the premises. Models may take the form of visual images, but their critical feature is their structure. The theory's fundamental principle is as follows:

The principle of truth: individuals tend to minimize the load on working memory by constructing mental models that represent explicitly only what is true, and not what is false.

The principle of truth is subtle, because it applies at two levels. First, individuals represent only true possibilities; second, they represent only those literal propositions in the premises—affirmative or negative—that are true in the true possibilities. For example, an inclusive disjunction such as

The flaw is in memory or in the software, or both
has three mental models that represent the true possibilities; that is, they correspond to the rows in a truth table

We thank the members of our laboratory, who have provided much helpful advice: Patricia Barres, Victoria Bell, Zachary Estes, Mary Newsome, Fabien Savary, Lisa Torreano, Isabelle Vadeboncoeur, and Yingrui Yang. We thank Geoffrey Keene for advice on modal logic, Lance Rips for testing his PSYCOP program with an illusory inference, Bonnie Meyer for carrying out the "think aloud" study, and Walter Schaecken and Robert Mackiewicz for a suggestion about formal rule theories. We also thank many other colleagues, too numerous to name, for their ideas and alternative hypotheses. We are grateful to Tom Ward, Simon Handley, and an anonymous reviewer, for their helpful criticisms. Correspondence should be addressed to Y. Goldvarg, Department of Psychology, Green Hall, Princeton University, Princeton, NJ 08544 (e-mail: goldvarg@princeton.edu).

that are true. We represent these three models in the following diagram:

Memory	
	Software
Memory	Software

where each row denotes a separate model, "memory" denotes a model of the flaw being in memory, "software" denotes a model of the flaw being in the software, and the third model combines these two models into one. If naive individuals are asked to list the possibilities given the disjunction above, they tend to list precisely the preceding models (Johnson-Laird & Savary, 1996). Each model represents only what is true in a particular possibility. Hence, for the first possibility, the model represents that it is true that the flaw is in memory. It is also false in this case that the flaw is in the software, but, according to the principle of truth, models do not usually represent false information explicitly. Similarly, for the second possibility, the model represents that the flaw is in the software, but it does not represent explicitly that it is false that the flaw is in memory. The theory postulates that reasoners make "mental footnotes" to keep track of the information about what is false, but that these footnotes are soon likely to be forgotten.

In contrast to mental models, *fully explicit* models of the disjunction represent the false components in each possibility, using negations that are true:

Memory	\neg Software
\neg Memory	Software
Memory	Software

where " \neg " denotes negation. These three fully explicit models match the three rows that are true in a truth table of the assertion. It is important to bear in mind that the principle of truth concerns the failure to represent what is false, not a failure to represent negation. If a negative as-

Table 1
The Sets of Mental Models for Each of the Major Sentential Connectives and Their Corresponding Fully Explicit Models

Connective	Mental Models		Fully Explicit Models	
A and B:	A	B	A	B
A or else B:	A		A	\neg B
		B	\neg A	B
A or B, or both:	A		A	\neg B
		B	\neg A	B
	A	B	A	B
If A then B:	A	B	A	B
		...	\neg A	B
			\neg A	\neg B
Iff A then B:	A	B	A	B
		...	\neg A	\neg B

Note—Each line represents a model of a separate possibility, “ \neg ” denotes negation, and “...” denotes a model with no explicit content. “Iff” denotes “if and only if,” and the only difference between its mental models and those for “if” are in the mental footnotes that are used to construct fully explicit models.

section in a proposition is true, then it will be represented, but cases where it is false will not be represented. Table 1 presents the mental models and the fully explicit models for each of the main sentential connectives, as generated by a computer program implementing the theory at different levels of expertise. Fully explicit models can be constructed from mental models provided that one has access to the footnotes indicating what is exhaustively represented in the mental models (Johnson-Laird & Byrne, 1991).

The model theory applies to modal reasoning—that is, reasoning about what is possible and what is necessary (Bell & Johnson-Laird, 1998). Other theories of reasoning might also be able to explain the phenomenon of modal reasoning. Thus, it might be feasible to complete Osherson’s (1976) formal rules for modal reasoning in such a way that the lengths of the derivations of the various sorts of conclusion predict their relative difficulty. Alternatively, current theories of necessary deductions (e.g., Braine & O’Brien, 1991, 1998; Rips, 1994) might be extended to deal with modal reasoning by adding formal rules for drawing conclusions about possibilities. Hence, a key question is whether the mental model theory predicts some unique phenomenon that would be difficult, if not impossible, for the other current theories of reasoning to predict.

The answer depends on the principle of truth. It predicts that naive human reasoners—that is, those with no training in logic—are, in effect, programmed to commit systematic fallacies, because of their failure to represent what is false. Consider, for instance, an exclusive disjunction of the form

p or else q or else r

in which only one of the three propositions (p, q, r) is true and the other two are false. According to the model theory, naive reasoners will construct the following sort of models for this disjunction:

P

Q

R

where P denotes the set of models corresponding to the proposition p, Q denotes the set of models corresponding to the proposition q, and R denotes the set of models corresponding to the proposition r. Hence, when individuals think about the truth of the proposition p, they should not bring to mind the falsity of the other two propositions, q and r, especially if the propositions p, q, and r are themselves of any degree of complexity. In most cases, reasoners can still reach the correct conclusion even if they fail to represent what is false, but in some cases, as we shall see, the theory predicts that they will draw an invalid conclusion.

Suppose that p, q, and r, are themselves disjunctive propositions, say, about a particular hand of cards:

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

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

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

We cannot express the disjunction of disjunctions with a premise of the form: “There is a king in the hand or there is an ace, or both, or else there is a queen in the hand or there is an ace, or both, or else, there is a jack in the hand or there is a ten, or both.” That would be too confusing. But, there is a simpler alternative way to express the disjunction of disjunctions, using the following rubric:

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.

This formulation is equivalent to the following: p or else q or else r, where these variables have as values the appropriate disjunctions above. The two ways of expressing the exclusive disjunction have the same truth conditions, the same mental models, and the same fully explicit models. Given the preceding premises and asked the following question:

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

reasoners who fail to represent what is false should succumb to an *illusion of possibility*. They should respond erroneously, “yes.” The principle of truth implies that when they consider the case that the first disjunction is true, they will fail to represent the falsity of the second and third disjunctions; when they consider the case that the second disjunction is true, they will fail to represent the falsity of the first and third disjunctions, and so on. The first disjunction has the following mental models (Table 1):

King

Ace

King Ace

They should respond "Yes, it is possible for there to be an ace in the hand," because two of the models of the possibilities contain an ace. The second disjunction also supports the same conclusion. Hence, reasoners should respond "yes." But, the conclusion is a fallacy.

When the first disjunction is true, the second disjunction is false:

\neg Queen \neg Ace

and the third disjunction is false:

\neg Jack \neg Ten

Hence, the fully explicit model of this case is

King \neg Queen \neg Ace \neg Jack \neg Ten

When the second disjunction is true, the first and third disjunctions are false, and the fully explicit models of this case are

\neg King Queen \neg Ace \neg Jack \neg Ten

And when the third disjunction is true, the first and second disjunctions are false, and the fully explicit models of this case are

\neg King \neg Queen \neg Ace Jack \neg Ten

\neg King \neg Queen \neg Ace \neg Jack Ten

\neg King \neg Queen \neg Ace Jack Ten

The set of these five fully explicit models of the premises as a whole show that in no case is there an ace present in the hand. Thus, when the falsity of the two remaining premises is taken into account in considering the potential truth of each premise, the correct response to the question emerges: "No, it is not possible for there to be an ace in the hand."

In other words, if there were an ace in the hand, then two of the premises would be true, contrary to the rubric that only one of them is true. But, according to the theory, when naive reasoners consider the models of one premise, they should overlook the consequences of the falsity of the other two premises. Reasoners consider the models of the truth of each premise in isolation from the other premises—it is as though these premises cease to exist for a moment, and so the mental models fail to represent anything at all about them. But, reasoners need to take their falsity into account in order to grasp the correct conclusion.

An analogous difficulty should arise with *illusions of impossibility*. These problems were created by replacing the two occurrences of "there is an ace" in the illusion above with "there is not an ace." Reasoners should now tend to respond wrongly, "no, it is not possible for there to be an ace in the hand."

A feasible research strategy would be to demonstrate that the fallacies occur systematically, just as one might demonstrate the occurrence of a visual illusion or an illusion in judging a probability (e.g., Tversky & Kahneman, 1973). Our aim, however, was to try to corroborate the model theory and to rule out some alternative explanations. One such explanation is that the problems are so complicated that they confuse naive reasoners, who, as a consequence, succumb to the illusions. (Many people have suggested such an account when we have presented them with an illusory inference informally.) A natural control for such an account is to use problems of a comparable complexity, but that should elicit a correct conclusion even if the participants fail to represent what is false. If reasoners err with the control problems, then the model theory is wrong and the cause of the errors may well be the complexity of the materials or the task. But, if the participants draw the correct conclusions to the control inferences and succumb to the illusory inferences, then the model theory is corroborated.

As a control for illusions of possibility, we used inferences of comparable syntactic complexity, but in which there was only one disjunction containing the clause "there is an ace." Participants should tend to respond correctly, "yes," even if they failed to represent what is false. As a control for illusions of impossibility, we needed an inference that was of comparable syntactic complexity, but to which reasoners should respond correctly, "no, an ace is not possible," even if they fail to represent the falsity of premises. If a problem has only a single occurrence of "there is not an ace" in one of the disjunctions, then an ace is still possible. In order to rule out its possibility, it is therefore necessary to use a *conjunctive* rubric:

Both of the following premises are true about a particular hand of cards:

There is a king in the hand and there is not an ace.

There is a queen in the hand and there is not an ace.

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

The mental models support the correct response, "no." Although conjunctions may be simpler syntactically than disjunctions, there is a compensating advantage of this sort of control problem with its conjunctive rubric. The use of two different rubrics in the experiment should help the participants to pay attention to them because they differ from one problem to another.

If naive reasoners rely on mental models to make modal inferences, then they should be susceptible to illusory inferences, which can be obviated by representing both the true and the false contingencies. Even if people follow the principle of truth, they should arrive at correct conclusions for most modal inferences, but there are some inferences where the failure to represent falsity should lead them systematically astray. The aim of our first experiment was to discover whether these predicted cognitive illusions occur. Our subsequent experiments pin down the source of the

phenomenon more precisely by eliminating some obvious alternative explanations.

EXPERIMENT 1

The experiment was designed to check whether illusions in modal reasoning occur when an exclusive disjunction is expressed in an explicit rubric:

Only one of the following premises is true:

and the propositions in the disjunction themselves are inclusive disjunctions or conditionals. The experiment also examined performance with control inferences to which the participants should draw correct conclusions even if they fail to represent those premises that are false when they consider the truth of a particular premise. The experiment accordingly examined four main sorts of inference:

1. Illusions of possibility, such as:

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?

Participants should tend to respond wrongly, "yes," to these problems.

2. Control inferences of possibility, which were of comparable syntactic complexity to the illusions of possibility, and which were constructed by having only one disjunction containing the clause "there is an ace." Participants should tend to respond correctly, "yes," to these problems.

3. Illusions of impossibility to which reasoners should respond wrongly, "no," and which were created by replacing the two occurrences of "there is an ace" in Problem 1 with "there is not an ace."

4. Control inferences of impossibility, which were of comparable syntactic complexity to the illusions of impossibility. If a problem has only a single occurrence of "there is not an ace" in one of the disjunctions, then an ace is still possible. In order to rule out its possibility, it is therefore necessary to use a *conjunctive* rubric:

Both of the following premises are true about a particular hand of cards:

There is a king in the hand and there is not an ace.

There is a queen in the hand and there is not an ace.

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

Reasoners should respond correctly, "no, an ace is not possible," to these problems even if they fail to represent the falsity of premises.

Method

Design and Materials. The participants carried out eight illusory problems and eight control problems, which were presented in

a different random order to each participant. Four of the illusory problems were designed to yield the conclusion that a card was possible (illusions of possibility), and four were designed to yield the illusion that a card was impossible (illusions of impossibility). The control problems were similarly divided into four problems in which a card was possible and four problems in which a card was impossible. The illusions were based either on disjunctions of disjunctions (as in the preceding examples) or on disjunctions of conditionals. The control problems were based on combinations of disjunctions, conditionals, or conjunctions. Table 2 presents the full set of problems, which are stated in an abbreviated form with the same hand of cards. In the experiment proper, each problem included the rubric, and the premises and question concerned different cards, so that no participant encountered the same hand of cards more than once. The participants chose their response from three possibilities: "yes," "no," "it can't be deduced from the premises." After each problem, they rated their confidence in their response on a 5-point scale (1 = *no confidence*; 5 = *complete confidence*).

Participants. Twenty Princeton undergraduate students participated in the experiment, which lasted for about 20 min. They received either a class credit or \$4 for their participation. None of them had received any training in logic or had participated before in any experiment on reasoning.

Procedure. The participants were tested individually in a quiet room. They read the instructions, which included the following remarks:

Please work on these problems carefully, and choose a solution for each of them. After you have made each choice, please mark on the 5-point scale how confident you are in your answer.

The instructions emphasized that it was important to pay attention to the rubric of each problem:

Though these problems may seem easy to you, some of them are very difficult. It is important to pay attention to the statements [i.e., the rubrics] before the premises. For some of the problems all of the premises are true, while for others only one premise is true and the rest are false. Please pay attention to the statement before each set of premises, and think carefully.

After the experimenter had answered any questions, the participants were given a simple practice problem. They were not allowed to use paper and pencil, but they could take as much time as they needed to answer each question in the experiment.

Results and Discussion

The overwhelming majority of the participants' responses were either "yes" or "no," with only 3% "can't be deduced from the premises," which we did not count as an error. Table 3 presents the overall percentages of correct responses and the mean confidence ratings. The percentages for the individual problems are shown in Table 2. The participants were correct on 91% of the control inferences, but only on 15% of the illusory inferences, and all 20 of the participants were more accurate with the control inferences than with the illusory ones ($p = .520$, i.e., less than 1 in a million). One unexpected finding, as Table 3 shows, was that the illusions of possibility were more compelling in comparison with their controls than the illusions of impossibility in comparison with their controls, and this interaction was reliable (Wilcoxon test, $z = 2.75$, $p < .006$). It is a result of the difference between the two sorts of illusions (Wilcoxon test, $z = 2.83$, $p < .005$), because there was no significant difference between their respective control problems (Wilcoxon test, $z = .61$). The participants were highly confident in all

Table 2
The Problems in Experiment 1 in an Abbreviated Form

Illusory Problems	% Correct	Control Problems	% Correct
Inferences of Possibilities			
If king then ace		If king then ace	
If queen then ace		If queen then jack	
If jack then ten		If jack then ten	
Is an ace possible?	5	Is an ace possible?	90
If king then ace		If king then ace	
If not a king then ace		If not a king then queen	
If jack then ten		If jack then ten	
Is an ace possible?	0	Is an ace possible?	95
King or ace or both		King or ace or both	
Queen or ace or both		Queen or jack or both	
Jack or ten or both		Ten or nine or both	
Is an ace possible?	0	Is an ace possible?	85
King or ace or both		King or ace or both	
Not a king or an ace or both		Not a king or jack or both	
Jack or ten or both		Queen or ten or both	
Is an ace possible?	0	Is an ace possible?	90
Inferences of Impossibilities			
		Both are true:	
If king then not an ace		If king then not an ace	
If queen then not an ace		If not a king then not an ace	
Is an ace possible?	30	Is an ace possible?	95
		Both are true:	
If king then not an ace		If queen then not an ace	
If not a king then not an ace		If not a queen then not an ace	
Is an ace possible?	30	Is an ace possible?	100
		Both are true:	
King or not an ace or both		King and not an ace	
Queen or not an ace or both		Queen and not an ace	
Is an ace possible?	30	Is an ace possible?	75
		Both are true:	
King or not an ace or both		King and not an ace	
Not a king or not an ace or both		Not a king and not an ace	
Is an ace possible?	30	Is an ace possible?	100

Note—Each problem was preceded by the rubric “Only one of the following premises is true about a particular hand of cards,” except for the control inferences of impossibility, which were preceded by “Both of the following premises are true about a particular hand of cards.” The questions concerned possibilities, and the table shows the percentages of correct answers.

their answers (Table 3), and there was no reliable difference between their ratings for the control problems (4.20 on the 5-point scale) and their ratings for the illusions (4.24 on the 5-point scale; Wilcoxon test, $z = .96$).

In general, the results corroborated the model theory. Nearly everyone succumbed to the illusions of possibility, and most participants succumbed to one or more illusions of impossibility. Hence, reasoners untrained in logic commit systematic fallacies in modal reasoning, for disjunctive assertions as well as conditional assertions, and where an exclusive disjunction of the assertions is expressed explicitly by a rubric of the form, “Only one of the following premises is true.”

What causes the illusions? The model theory explains them in terms of an incomplete understanding of exclusive disjunctions, whether they are expressed by way of the rubrics in the present experiments:

Only one of the following premises is true:

or by the sentential connective *or else*, which gave rise to illusions in an experiment on deductively necessary conclusions (Johnson-Laird & Savary, in press). In both cases, according to the theory, reasoners consider the truth of each premise, but when they do so, they fail to represent the falsity of the other premises. Theories based on formal rules postulate distinct processes of understanding and reasoning, but theories based on mental models make much less of a distinction. Thus, Polk and Newell (1995, p. 563) wrote that their theory explains reasoning in terms of linguistic processes “without the need to posit reasoning-specific mechanisms.” We do not go quite that far, but the present theory does postulate that the cause of the illusions is a failure to construct a completely explicit representation of exclusive disjunctions.

Could it be that the cause of the illusions is merely that the participants paid insufficient attention to the rubrics, or indeed forgot about them? We are confident that such a lack of attention is not the cause of the illusions, because

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

Inference Type	Illusions		Controls	
Possibility	1	(4.17)	90	(4.13)
Impossibility	29	(4.23)	92	(4.35)

Note—The figures in parentheses represent the participants' mean confidence in their answers on a 5-point scale (1 = *no confidence*; 5 = *complete confidence*).

many participants spontaneously remarked that any *single* premise could be true for the illusions of possibility, for example, "So I can choose any of the premises to be true. If I choose the first or the second one, an ace is there." Another remark about such a problem was, "Yes, an ace is possible if the first or the second premise is true." (We have changed the name of the card in these protocols to fit Problem 1 above.) What such protocols show is that the participants grasped that the premises were in a disjunction, but treated a disjunction as though it meant merely that any one premise could be true, overlooking that when one premise is true, the other premises are false. We will return to the possibility that the participants ignored the rubric in the General Discussion.

In order to use control problems to which the correct conclusion was that a card was impossible, it was necessary to use conjunctive rubrics, which asserted that both premises were true. The syntactic complexity of the corresponding illusions might have been greater and somehow misled the participants in ways that did not occur with the control problems. Such an explanation is implausible if only because the illusions of impossibility were less powerful than the illusions of possibility. In general, however, the illusions were based on different premises from the control problems, and so our second experiment was designed to eliminate the difference. Likewise, the rubric "Only one of the following premises is true" might have encouraged the participants to neglect falsity, and so the rubrics in the second experiment explicitly mentioned both truth and falsity.

EXPERIMENT 2

In this experiment, we used the *same* premises for four sorts of inferences: illusions of possibility, their controls with the correct answer "yes," illusions of impossibility, and their controls with the correct answer "no." In the previous experiment, the illusory inferences were based on the rubric "Only one of the following premises is true," and so the participants had to infer that the other premises were false. This inference might have contributed to the difficulty of the illusory inferences. Hence, in the present experiment, we used a rubric that was explicit about both truth and falsity, as the following problem shows:

One of the following premises is true about a particular hand of cards and one is false:

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

There is a queen in the hand and there is an ace.

These premises were combined on four separate trials with questions of the following form (except that in the experiment proper each set of premises and its single question used a different hand of cards):

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

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

Is it possible that there is a queen in the hand and not an ace?

Is it possible that there is not a king in the hand and not an ace?

Once again, the model theory predicts that when naive reasoners consider the truth of one of the premises, they will fail to take into account the concurrent falsity of the other premise. When they consider the truth of the first premise, they will envisage the following mental models of its true possibilities:

K	
	A
K	A

When they consider the truth of the second premise, which is a conjunction, they will envisage the mental model of the one true possibility:

Q	A
---	---

These models, of course, fail to take into account the falsity of one premise when the other premise is true. The fully explicit models for the premises as a whole, taking falsity into account, are as follows:

K	¬Q	A
K	Q	¬A
K	¬Q	¬A
¬K	¬Q	A

The premises should therefore yield an illusion of possibility to the first question (queen and ace?); that is, the participants should respond "yes" on the basis of their model of the second premise, but the fully explicit models show that the answer is wrong because whenever there is a queen, there is not an ace and vice versa. Indeed, if there were a queen and an ace in the hand, then both premises would be true. The second question (king and ace?) should also elicit the answer "yes," but this response is supported by the fully explicit models, and so it is correct. Likewise, the third question (queen and not an ace?) should elicit an illusion of impossibility, and the fourth question (not a king and not an ace?) should elicit a correct "no" answer.

Method

Design and Materials. The participants carried out eight illusory problems (four illusions of possibility and four illusions of impossibility) and their eight respective control problems. The prob-

Table 4
The Problems in Experiment 2 in an Abbreviated Form,
Their Mental and Their Fully Explicit Models

Form of Premises and Questions	Mental Models		Fully Explicit Models			%Correct
1. If king then ace If queen then ace	K	A	$\neg K$	Q	$\neg A$	
		Q	A	K	$\neg Q$	$\neg A$
	...					
king and ace?	(Illusion of possibility)					5
king and not queen?	(Control "yes" response)					75
king and not ace?	(Illusion of impossibility)					25
not queen and not king?	(Control "no" response)					80
2. If king then not ace If queen then not ace	K	$\neg A$	$\neg K$	Q	A	
		Q	$\neg A$	K	$\neg Q$	A
	...					
king and not ace?	(Illusion of possibility)					0
queen and not king?	(Control "yes" response)					95
king and ace?	(Illusion of impossibility)					50
not king and not queen?	(Control "no" response)					80
3. King or ace, or both Queen and ace	K		K	$\neg Q$	A	
		A	K	Q	$\neg A$	
	K	A	K	$\neg Q$	$\neg A$	
	Q	A	$\neg K$	$\neg Q$	A	
queen and ace?	(Illusion of possibility)					30
king and ace?	(Control "yes" response)					95
queen and not ace?	(Illusion of impossibility)					40
not king and not ace?	(Control "no" response)					95
4. Not king or not ace, or neither	$\neg K$		$\neg K$	Q	A	
		$\neg A$	$\neg K$	$\neg Q$	A	
Not queen and not ace	$\neg K$	$\neg A$	K	Q	$\neg A$	
	$\neg Q$	$\neg A$	$\neg K$	Q	$\neg A$	
not queen and not ace?	(Illusion of possibility)					15
not king and not ace?	(Control "yes" response)					75
not a queen and an ace?	(Illusion of impossibility)					55
king and ace?	(Control "no" response)					90

Note—Each pair of premises was prefaced by "One of the following premises is true and one is false." The questions concerned possibilities, and the table shows the percentages of correct answers. For convenience, we have stated each problem and question with the same sets of cards; in the experiment, the premises were presented with a single question on separate trials with different cards.

lems were presented in a different random order to each participant. All the problems included the rubric, and each set of premises and question concerned a different hand of cards, so that no participant encountered the same hand in more than one problem. Table 4 presents the set of problems in an abbreviated form, their mental models, and their fully explicit models. As in the previous experiment, the participants chose their response from three possibilities: "yes," "no," "it can't be deduced from the premises"; and they rated their confidence in their responses on a 5-point scale.

Participants. Twenty Princeton undergraduate students participated in the experiment, which lasted for about 20 min. They received a class credit for their participation. None of them had received any training in logic or had participated before in any experiment on reasoning.

Procedure. The participants were tested individually in a quiet room. They were given the same general instructions and practice problem as in Experiment 1. Once again, the instructions emphasized the importance of paying attention to the rubric of the problems: The participants were told that in every case, one of the two premises was true and one of them was false, though they would not know which was which. One other modification, however, was that

the participants had a choice of only two answers—"yes" or "no"—to each question. We dropped the option that a response could not be deduced from the premises, because it had been so rarely chosen in Experiment 1. As before, the participants were not allowed to use paper and pencil, but they could take as much time as they needed to answer each question.

Results and Discussion

Table 5 presents the overall percentages of correct responses and the mean confidence ratings. The percentages for the individual problems are shown in Table 4. The participants were correct on 79% of the control problems, but only on 22% of the illusory problems. All but 1 of the 20 participants were more accurate with the control problems than with the illusory ones (Sign test, $p < .0005$). As in Experiment 1, the illusions of possibility were more compelling in comparison with their controls than were the illusions of impossibility in comparison with their controls, and this interaction was reliable (Wilcoxon

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

Inference Type	Illusory Problems		Control Problems	
Possibility	8	(4.41)	85	(4.25)
Impossibility	36	(3.99)	73	(4.23)

Note—The figures in parentheses represent the participants' mean confidence in their answers on a 5-point scale (1 = *no confidence*; 5 = *complete confidence*).

test, $z = 2.82, p < .005$). The interaction was mainly a result of the difference between the two sorts of illusions (Wilcoxon test, $z = 3.025, p < .003$), because there was no significant difference between their respective control problems (Wilcoxon test, $z = 1.73$, n.s., two-tailed test). The participants were highly confident in all their answers (Table 5), and there was no reliable difference between their ratings for the control problems (4.24 on the 5-point scale) and their ratings for the illusions (4.20 on the 5-point scale; Wilcoxon test, $z = .12$).

The experiment confirmed the occurrence of illusions in modal reasoning. Because the same premises were used for both the illusions and controls, the results show that illusions are not a result of more complex or confusing premises than those of the control problems. All that differed between illusions and controls were the particular questions, which were all in the form of conjunctions (with or without negative assertions). As Table 4 shows, there is no obvious difference between the questions eliciting illusions and those eliciting control responses. Thus, for example, the question of the form

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

occurred for all four sorts of problem (the two sorts of illusion and their respective controls). Likewise, the participants were again highly confident in their responses to both control and illusory problems. We conclude that whatever the cause of the illusions, they cannot be attributed to some difficulty in understanding the materials, the task, or the instructions. Once again, the illusions of possibility were more compelling than the illusions of impossibility, and we will return to this phenomenon in the General Discussion.

EXPERIMENT 3

How persistent are the illusions? Is it possible for people to adopt a strategy that would help them to avoid the illusions? If they result from a neglect of what is false, then one antidote should be to emphasize the importance of falsity for reasoning. Several independent studies have shown that such instructions do reduce participants' susceptibility to illusions (e.g., Newsome & Johnson-Laird, 1996). Likewise, Tabossi, Bell, and Johnson-Laird (in press) used the following rubric:

Only one of the following premises is *false*

which presumably biased the participants to consider the consequences of the falsity of the respective premises. The result was a reliable reduction in illusions about probabilities, though performance rose to only 31% correct. One reason for the relatively small improvement may be that people do not have direct access to the conditions that would falsify assertions based on sentential connectives. They have to derive them from the conditions in which the assertions would be true. In deriving these false conditions, individuals make many errors (Barres & Johnson-Laird, 1997). Thus, even when the participants were told to think about falsity, they were likely to have made mistakes, or they may not have been able to avoid representing true models, and these models may have led them into illusions.

Is there some other antidote to the illusions that would not require reasoners to change their representation of the premises? Given the rubric "Only one of the premises is true," there is a way to try to reduce the illusions by invoking a simple additional step. If reasoners check the consequences of their conclusion for the truth of the rubric, then they should be able to determine whether their conclusion renders *more than one* conclusion true. This procedure has the advantage that reasoners do not even have to try the difficult task of envisaging the circumstances in which the premises would be false. They have merely to check how many premises their conclusion renders true. Experiment 3 investigated whether those participants who were told to make this additional check would perform better than those who were not given the instruction.

Method

Design and Materials. The participants were allocated at random to two independent groups, which each carried out 16 inferences (the eight illusions and eight control problems of Experiment 1) in a random order. Halfway through the experiment, participants in the "help" group were instructed to check whether their conclusion was consistent with the rubric that one premise was true and the rest were false, whereas the "No help" group received no such instruction. The materials were identical to those of Experiment 1.

Participants. Twenty Princeton undergraduate students participated in the experiment, which lasted for about 30 min. They received a class credit for their participation. None of them had received any training in logic or participated in any previous experiment on reasoning.

Procedure. The participants were tested individually in a quiet room. They received the instructions of Experiment 1 with a modification for the participants in the help group, who were given additional instructions after they had carried out half the problems:

Some of these problems are extremely tricky because only one of their premises is true. To solve these problems correctly, you need to do the following: (1) select your solution; (2) go back and check whether your solution preserves the relationship between the premises, that is, one of them is still true and the rest are false. For example:

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

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

If there is a queen in the hand then there is a jack in the hand.

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

Table 6
The Percentages of Correct Responses in Experiment 3

Group	Inference Type	Illusory Problems		Control Problems	
		First Half	Second Half	First Half	Second Half
No help	Possibility	0	0	100	100
	Impossibility	47	60	95	95
Help	Possibility	0	57	93	100
	Impossibility	70	85	90	87

Note—For the second half of the experiment, participants in the help group were told to check their conclusions against the rubric that only one of the premises was true.

If your first choice is “yes” (you believe that an ace is possible), then go back and check if one of the premises is still true and the other could still be made false. The first premise is definitely true since you have an ace. The second one is false as long as you have a queen and do not have a jack. Thus your choice is correct and an ace is possible. Please do this checking for every problem; without it you will get most of the problems wrong.

The participants in the no help group received no such instructions, but they had a comparable pause that allowed them to rest between the two halves of the experiment. Both groups, of course, received the instructions to pay attention to the rubrics, which differed from one problem to another, as in Experiment 1.

Results

Table 6 presents the two groups' percentages of correct responses to the four sorts of problems in the two halves of the experiment. Overall, the participants were correct on 95% of the control inferences, but only on 39% of the illusory inferences (Wilcoxon test, $z = 7.3$, $p < .001$). As in the two previous experiments, the illusions of possibility were more compelling in comparison with their controls than the illusions of impossibility in comparison with their controls, and this interaction was reliable (Wilcoxon test, $z = 2.82$, $p < .005$). There was no reliable difference between the confidence ratings for the control problems (mean confidence of 4.24) and those for the illusory problems (mean confidence of 4.17; Wilcoxon test, $z = .56$, $p > .05$).

The two groups did not differ in the first half of the experiment: The help group made 60% correct responses and the no help group made 63% correct responses (Mann-Whitney $U = 38.0$, $p > .3$). But, the effect of the instructions to the help group halfway through the experiment was exactly as predicted: They performed better than the no help group (Mann-Whitney $U = 10$, $p < .002$), and their improvement on the illusory problems in comparison with the controls was reliably greater than that of the no help group (Mann-Whitney $U = 8.0$, $p < .001$). Indeed, in the second half of the experiment, the help group made 71% correct responses to the illusions, whereas the no help group made only 30% correct responses.

The experiment corroborated the occurrence of illusory inferences, but it showed that a simple manipulation made the participants in the help group less susceptible to them. They performed reliably better after they had been

told to check that their responses were consistent with the rubric that only one of the premises was true. This additional check allowed the participants to notice that their initial answer made two of the premises true. All the participants who were then able to solve an illusory problem had succumbed to the illusion when they first tackled the problem. The check enabled them to correct their error. Even with this instruction, however, their performance was still far from perfect.

GENERAL DISCUSSION

Reasoning about possibilities has seldom been studied by psychologists. The mental model theory postulates that models of premises represent possibilities in much the same way that the true rows in a truth table represent possibilities (Johnson-Laird & Byrne, 1991). Possibilities are therefore at the heart of the theory. The principle of truth, however, specifies that mental models represent neither false possibilities nor the falsity of literal propositions in the premises within the models of true possibilities. Certain modal inferences should therefore be illusory; that is, they should have compelling, but invalid, conclusions. We have discovered that such illusions occur and that they are robust. Consider, for example, the following problem from Experiment 1:

1. 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?

Almost every participant answered “yes,” though the conclusion is a fallacy. The model theory predicted the illusion on the grounds that naive reasoners would consider the truth of each premise without taking into account the concomitant falsity of the other two premises. Thus, when they considered the truth of the first premise, they would infer that an ace was possible, overlooking that this possibility is ruled out by the concomitant falsity of the second premise. Indeed, if there is an ace in the

hand, then the first *two* premises would be true, contrary to the rubric that only one premise is true. The premises, however, are quite complicated, and so in order to try to pinpoint the cause of the illusions, we ran control problems to which reasoners should reach the correct conclusions even if they fail to represent what is false. Our aim was to match the illusory problems with control problems of comparable syntactic complexity. Thus, when only one of the premises above contained the clause, "there is an ace," the majority of participants correctly answered, "yes."

The results of Experiment 1 corroborated the model theory. They showed that illusions occur with disjunctions as well as with conditional premises, and that they occur when an exclusive disjunction between the premises is expressed explicitly using the rubric "only one of the following premises is true." The experiment used another sort of rubric for the control problems for illusions of impossibility: "Both of the following premises are true." This use of two rubrics should have helped participants to keep the rubrics in mind, but the illusions of impossibility might have been harder than their controls because the illusions used a disjunction of premises, whereas the controls used a conjunction of premises. In fact, the results showed no such effect: The illusions of impossibility were less compelling than the illusions of possibility, and there was no difference between their respective controls.

Nevertheless, a criticism of Experiment 1 is that the illusory premises were different from the control premises, and so the participants might have had, say, a greater difficulty in coping with the syntax of the illusory premises. We therefore carried out Experiment 2 using the *same* premises for both the illusions and the controls. Given the following problem:

2. One of the following premises is true and one is false about a particular hand of cards:

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

There is a queen in the hand and there is an ace.

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

most participants responded "yes," even though the presence of a queen and an ace renders *both* premises true. When the same premises were combined with a control question of the form

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

the majority of participants, as predicted by the model theory, correctly responded "yes." They were probably right for the wrong reason; that is, they still did not consider the consequences of falsity, but this oversight does not lead them into error with the control problems.

With hindsight, it is astonishing that 99% of the responses to Problem 1 above were "yes," because it seems obvious that the presence of an ace renders two of the premises true. A rubric that emphasizes falsity reliably reduces illusory inferences (Tabossi et al., in press), but it

does not eliminate them. The reason, we suspect, is that people have difficulty in grasping the conditions in which premises containing connectives are false. They have no direct access to these conditions, but have to infer them and they make many errors (Barres & Johnson-Laird, 1997). We therefore asked one group of participants in Experiment 3 to check whether their conclusions met the constraint that only one of the premises was true. This procedure obviates the need to envisage the falsity of the premises, and, as predicted, it led to a considerable decline in the illusions.

One unexpected phenomenon in all three of our experiments was that illusions of possibility were more compelling than illusions of impossibility. In Experiments 1 and 3, the illusions of possibility depended on different premises than the illusions of impossibility, but no explanation based on this difference can account for the results of Experiment 2, in which both sorts of illusion depended on the same premises. The most likely explanation derives from another prediction of the model theory of modal reasoning, which has been corroborated by Bell and Johnson-Laird (1998). To infer that a situation is possible is easier than to establish that it is impossible, because possibility depends on truth in only one model of the premises, whereas impossibility depends on falsity in all the models. It follows that reasoners will respond that a situation is possible as soon as they find a model that satisfies it, but that they will need to do much more work—and perhaps try to flesh out their models more explicitly—before they respond that a situation is not possible. This pattern should apply to both illusory and control problems. Indeed, it is apparent for the control problems in Experiments 2 and 3, though not for those in Experiment 1, which are closer to ceiling.

Let us try to be as clear as possible about what the model theory postulates. It predicts that people are susceptible to illusions because they do not fully understand sentential connectives, such as exclusive disjunction. That is, they think about the truth of each premise in the disjunction, but when they do so, they fail to take into account the concurrent falsity of the other premise. Our results corroborate this claim both for disjunctions of disjunctions and for disjunctions of conditionals. Another study of necessary conclusions has corroborated the claim when the principal disjunction is expressed, not by a rubric, but by the connective "or else" (Johnson-Laird & Savary, in press). The effect is not unique to cases in which the main connective is an exclusive disjunction, but occurs in probabilistic reasoning with cases in which the main connective is a biconditional and even a conjunction (Johnson-Laird & Savary, 1996).

Is there a plausible alternative explanation for the phenomenon? One possibility is that the problems, materials, or task are too difficult for individuals untrained in logic to understand. Our results count against this possibility. The participants were highly confident in their responses, whether illusory or correct, which shows that

they were not consciously perplexed by the materials or task. Moreover, we used the *same* premises for both the illusions and controls in Experiment 2. The two sorts of problems differed only in the questions that were put to the participants. These questions were always conjunctions, and the particular questions were rotated from one sort of problem to another.

Another possible explanation of our results is that the participants failed to notice the rubrics, forgot about them, or mistakenly interpreted them in some other way. In other words, they failed to realize that

Only one of the following premises is true (Experiments 1 and 3)

and

One of the premises is true and one is false (Experiment 2)

mean that only one of the premises can be true, and treated them instead as though they meant that *both* premises were true. Several results run counter to this explanation. First, as we reported in the results of Experiment 1, the participants' protocols showed that they did indeed interpret the disjunctive rubrics as disjunctions. They remarked spontaneously that the truth of the first or the second premise in Problem 1 above was sufficient for there to be an ace. None of the protocols implied that a participant had made a conjunctive interpretation of the premises (see also Newsome & Johnson-Laird, 1996; Tabossi et al., in press). Second, in collaboration with Bonnie Meyer, we carried out a study using "or else" in which the participants had to think aloud as they responded to the following sort of problem (and others):

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

if there isn't an ace in the hand then there is a six in the hand.

There is an ace in the hand.

What, if anything, follows?

All the participants inferred that there was a six in the hand. Yet, as we explained earlier, it is a fallacy. A typical "think aloud" protocol from a participant warned that the problems are tricky is as follows:

If there is an ace there is a six or else if there is not an ace there is a six. So there's always a six. There is an ace. . . . If there is an ace then there is a six or if there is not an ace then there is a six. So, either way you have a six, and you have an ace. So, you have an ace and a six—I guess. If there is an ace there is a six . . . or if there is not an ace then there is a six. Yes. So regardless you have a six. I think. . . . So there has to be something else because that just seems way too easy. And something is wrong somewhere—I don't know what it is.

Much of the protocol consists in the participant repeating the premises, making the invalid inference, and wondering whether something may have gone wrong. The phrase "either way you have a six" shows that the par-

ticipant clearly had made a disjunction of the two conditionals, not a conjunction of them. Like this protocol, the others made clear that the participants interpreted the premises as disjunctions. Third, a study of illusions in probabilistic reasoning (Johnson-Laird & Savary, 1996) made use of three sorts of rubric within the same experiment: exclusive disjunctions, conjunctions, and biconditionals. In each case, the predicted illusions occurred, and performance on corresponding control problems was good, which would have been impossible if the participants had paid no attention to the rubrics. Finally, in Experiment 3, even the participants in the group that were instructed halfway through the experiment to check that their conclusions were consistent with one premise being true and the others being false, and that, as predicted, were less susceptible to the illusions, still performed much worse with illusions of possibility (57% correct) than with the control problems (100% correct). Indeed, we—and other researchers—have explored a variety of remedial procedures, but no one has so far discovered a manipulation that raises performance on illusory problems up to nearly flawless. The task would be trivial if it were merely a matter of ensuring that individuals pay attention to the rubrics of problems.

Of course, no empirical observations can rule out the possibility that some individuals sometimes fail to understand the task, sometimes adopt an idiosyncratic interpretation of a premise, and sometimes ignore a rubric. None of these alternatives, however, can explain all the illusions in reasoning about what is possible, necessary, and probable. Yet, they were all predicted by the principle that people cannot cope with falsity. In contrast, if individuals typically represented falsity, they would have avoided the illusions.

The illusory inferences are not predicted by any current theory of reasoning based on formal rules of inference (see Braine & O'Brien, 1991; Rips, 1994). Could these theories or Osherson's (1976) account of modal reasoning be modified so that they would explain the illusions? The short answer is that no one knows. But, what we can say is that the modification will not be easy. The illusions are systematically invalid conclusions, and current theories of reasoning based on formal rules make use only of rules that deliver *valid* conclusions. They have no machinery, as yet, to generate invalid conclusions in a systematic way. One tempting way to try to salvage formal rules is, first, to construe the model theory as postulating the rule—when people think about one disjunct they forget the others—and, second, to argue that if the model theory can have such a rule, then so can theories based on formal rules. In fact, the mental model theory has *no* such rule. It assumes that people represent only what is true. This principle has the effect that when people think about one *disjunct* it is as though they cease to think about the other. But, it is important to realize that the principle of truth underlies what is going on. If people merely ceased to think about one proposition and then about another, their performance would not be affected by different rubrics,

such as biconditional ones, or by the equivalent sentential connectives. Indeed, not all illusory inferences depend on exclusive disjunctions. As we mentioned earlier, there are compelling illusions in probabilistic reasoning based on biconditional rubrics (Johnson-Laird & Savary, 1996).

A more plausible way to salvage formal rule theories is to argue that reasoners somehow misapply a suppositional strategy (L. Bonatti, personal communication, April 29, 1997). Current formal rule theories have no principles that allow for such a misuse of suppositions or for inferring possibilities. It may be possible to frame the required principles, but such a misuse of suppositions appears to break down with the certain problems. Consider a problem from Experiment 2 based on a disjunction of conditionals:

One of the following assertions is true and one of them is false:

If there is a king then there is an ace.

If there is a queen then there is an ace.

4. Is it possible that there is a king and an ace?

Reasoners could consider the first premise, and then make the supposition

There is a king.

They could infer that there is an ace, and so succumb to the illusory inference that it is possible that there is a king and an ace. But, now consider a control problem based on the same premises, but with the following question:

5. Is it possible that there is a king and not a queen?

There is no way that a supposition of a king yields the conclusion that there is not a queen, and so reasoners should respond "no" to the question. In fact, the majority of participants *correctly* responded "yes." In other words, a misapplication of suppositions here obliterates the distinction between illusory and control problems (Yang & Johnson-Laird, in press). It remains an open question whether any simple modification to current formal rule theories can account for the phenomenon.

The difficulty in coping with falsity accounts for many of the problems that arise in reasoning. It explains the difficulty of modus tollens inferences (see Evans, Newstead, & Byrne, 1993). It also explains the difficulty of Wason's selection task in which participants are asked which of four cards,

A B 2 3

they need to turn over in order to find out whether a conditional assertion is true or false:

If there is an A on one side of a card then there is a two on the other side of the card.

The participants already know that each card has a letter on one side and a number on the other side. They typically err by failing to select the three card (see, e.g., Wason &

Johnson-Laird, 1972). If they fail because they do not represent the cases in which the conditional would be false, then any manipulation that helps them to take into account what is false should improve performance, and this hypothesis has been widely corroborated (see, e.g., Green & Larking, 1995; Liberman & Klar, 1996; Love & Kessler, 1995; Sperber, Cara, & Girotto, 1995).

Illusory inferences do occur in daily life, but they are so compelling that they probably go unnoticed. A search of the World-Wide Web turned up several cases. For example, a professor warns students absent from his class:

... either a grade of zero will be recorded if your absence is not excused, or else if your absence is excused other work you do in the course will count.

What he has in mind are clearly two possibilities, which are represented by the mental models of his assertion:

\neg excused	zero-grade
excused	other-work-counts

But, these models are illusory. The connective "or else" means that one assertion is true and one is false, but what the professor means is that *both* of his assertions are true. To convey the possibilities correctly, he should not have used "or else," but rather "but," which has the force of a conjunction:

... a grade of zero will be recorded if your absence is not excused, but if your absence is excused other work you do in the course will count.

His actual warning asserts an unintended set of possibilities, which allow, for instance, that students may not get a zero grade even when they have no excuse for their absence. Such errors show the pervasiveness of the illusions, which in turn gives us an insight into how human beings reason. In an age of increasing emphasis on the logic of technology, it is only a matter of time before the failure to consider falsity leads to errors with more serious consequences.

We conclude that individuals untrained in logic rely on mental models to make modal inferences. These models represent only what is true, and so reasoners commit systematic fallacies in cases where it is crucial to take falsity into account. To rely on truth is a sensible way to deal with limited processing capacity. In some cases, however, it leads reasoners into the dangerous illusion that they grasp a possibility that is in fact beyond them. The simple procedure that we used as an antidote to illusions is not always applicable, and so there is a need to develop other antidotes.

REFERENCES

- BARRES, P. E., & JOHNSON-LAIRD, P. N. (1997). Why is it hard to imagine what is false? *Proceedings of the Nineteenth Annual Conference of the Cognitive Science Society* (p. 859). Mahwah, NJ: Erlbaum.
- BELL, V., & JOHNSON-LAIRD, P. N. (1998). A model theory of modal reasoning. *Cognitive Science*, 22, 25-51.

- 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.
- BRAINE, M. D. S., & O'BRIEN, D. P. (Eds.) (1998). *Mental logic*. Mahwah, NJ: Erlbaum.
- EVANS, J. S. B. T., NEWSTEAD, S. E., & BYRNE, R. M. J. (1993). *Human reasoning: The psychology of deduction*. Hillsdale, NJ: Erlbaum.
- GREEN, D. W., & LARKING, R. (1995). The locus of facilitation in the abstract selection task. *Thinking & Reasoning*, **1**, 183-199.
- JOHNSON-LAIRD, P. N., & BYRNE, R. M. J. (1991). *Deduction*. Hillsdale, NJ: Erlbaum.
- JOHNSON-LAIRD, P. N., & SAVARY, F. (1996). Illusory inferences about probabilities. *Acta Psychologica*, **93**, 69-90.
- JOHNSON-LAIRD, P. N., & SAVARY, F. (in press). Truth and illusion in reasoning. *Cognition*.
- LIBERMAN, N., & KLAR, Y. (1996). Hypothesis testing in Wason's selection task: Social exchange cheating detection or task understanding. *Cognition*, **58**, 127-156.
- LOVE, R., & KESSLER, C. (1995). Focussing in Wason's selection task: Content and instruction effects. *Thinking & Reasoning*, **1**, 153-182.
- NEWSOME, M. R., & JOHNSON-LAIRD, P. N. (1996). An antidote to illusory inferences? In G. W. Cottrell (Ed.), *Proceedings of the Eighteenth Annual Conference of the Cognitive Science Society* (p. 830). Mahwah, NJ: Erlbaum.
- OSHERSON, D. N. (1976). *Logical abilities in children. Vol. 4: Reasoning and concepts*. Hillsdale, NJ: Erlbaum.
- 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.
- SPERBER, D., CARA, F., & GIROTTO, V. (1995). Relevance theory explains the selection task. *Cognition*, **52**, 3-39.
- TABOSI, P., BELL, V. A., & JOHNSON-LAIRD, P. N. (in press). Mental models in deductive, modal, and probabilistic reasoning. In C. Habel & G. Rickheit (Eds.), *Mental models in discourse processing and reasoning*. Berlin: John Benjamins.
- TVERSKY, A., & KAHNEMAN, D. (1973). Availability: A heuristic for judging frequency and probability. *Cognitive Psychology*, **4**, 207-232.
- WASON, P. C., & JOHNSON-LAIRD, P. N. (1972). *The psychology of deduction: Structure and content*. Cambridge, MA: Harvard University Press.
- YANG, Y., & JOHNSON-LAIRD, P. N. (in press). Illusions in quantified reasoning: How to make the impossible seem possible and vice versa. *Memory & Cognition*.

(Manuscript received August 11, 1997;
revision accepted for publication January 14, 1999.)