Cognition, 36 (1990) 69–84 **Discussion**

Meta-logical problems: Knights, knaves, and Rips*

P.N. JOHNSON-LAIRD
Princeton University

RUTH M.J. BYRNE
University of Wales College of Cardiff

Johnson-Laird, P.N., and Byrne, R.M.J., 1990. Meta-logical problems: Knights, knaves, and Rips. Cognition, 36: 69–84.

Introduction

In a pioneering study, Rips (1989) reports an investigation of what we will refer to as "meta-logical" puzzles. The puzzles he studied depend on imagining that there are only two sorts of persons: knights, who always tell the truth; and knaves, who always lie. Here is a typical problem: "There are two inhabitants, A and B, each of whom is a knight or a knave. A says, 'I am a knave and B is a knave'. B says, 'A is a knave'. What is the status of A and B: knight, knave, or impossible to tell?" The solution is that A is a knave and B is a knight.

Rips makes two principal claims. First, he argues that the process of solving these problems is accurately characterized by his theory, which uses *formal* rules of inference in the natural deduction format, and which is an extension of an earlier theory (see Rips, 1983). Second, he presents a challenge to theorists who espouse mental models: "Produce an explicit account of reasoning on knight–knave problems that is (a) theoretically explicit, (b) empirically adequate, and (c) not merely a notational variant of the natural-deduction theory" (Rips, 1989, p. 113). Our aim in this paper is to meet this challenge. We will present an alternative theory based on mental models, which is not a notational variant of the natural deduction theory, which is explicit, and which provides a better theoretical account of Rips's results.

*Requests for reprints should be sent to P.N. Johnson-Laird, Department of Psychology, Green Hall, Princeton University, Princeton, NJ 08544, U.S.A.

A critique of Rips's theory

We have four main misgivings about Rips's account. First, he underestimates the importance of meta-logical thinking. He appears to view it as logically on a par with other sorts of deduction. He overlooks a key feature: it hinges on an explicit concern with matters of truth and falsity. There is an irony here. The heyday of the purely formal approach to logic (as opposed to psychology) was brought to an end by the development of meta-logic: Tarski showed how to formulate a model-theoretic semantics for the syntactic calculus of predicate logic, and Gödel proved that the predicate calculus is "complete" in that it provides a formal derivation for any theorem that is valid in the semantics (e.g., Boolos & Jeffrey. 1980). Yet, Rips himself eliminates "truth" and "falsity" in his theory. Elsewhere, he has written: "cognitive psychology has to do without semantic notions like truth and reference that depend on the relationship between mental representations and the world" (Rips, 1986). As a general analysis of meta-linguistic assertions, the problems of this approach are well known (see Austin, 1970; Barwise & Etchemendy, 1987). Without the notion of truth, there is no notion of validity, and meta-logic evaporates.

Second, Rips deals with just one variety of meta-logical puzzle. There are many other varieties, for example: "There are two sorts of persons: logicians, who always make valid deductions; and politicians, who never make valid deductions. A says that either B is telling the truth or else B is a politician (but not both). B says that A is lying. C deduces that B is a politician. Is C a logician?" Rips's program cannot handle all knight-and-knave problems—not even all of those that he investigated experimentally, and it cannot handle other sorts of problem without introducing a vast number of formal rules.

Third, Rips proposes only a single deterministic procedure for solving meta-logical problems. According to this procedure, reasoners assume that the first individual in a puzzle is a knight, and explore the consequences of this assumption; next, they assume that this individual is a knave, and explore the consequences of this assumption. They then assume that the next individual in a puzzle is a knight, and so on. In fact, we have observed that logically untutored individuals are much less systematic in their approach until they have had considerable experience with the problems. They do not come to the task armed with a deterministic procedure that leads to the solution. Unlike, say, a linear syllogism, where the answer emerges rapidly and almost automatically, people can and do reflect about these problems. They spontaneously make meta-logical remarks, for example: "A could be telling the truth about the first part [of the conjunction] so the second part of A's assertion is a lie". In short, an adequate theory must allow for diverse strategies in solving, or attempting to solve, meta-logical problems. The procedure embodied in Rips's program is too inflexible.

Fourth, and most important, the procedure postulated by Rips is too powerful and, as we will show in the next section, it places an impossible load on the processing capacity of working memory.

A theory of meta-logical reasoning based on mental models

Reasoning with models

In our view, the ability to make straightforward propositional deductions is a prerequisite for solving meta-logical puzzles. For example, reasoners must be able to make the following deduction:

A or B (or both)
not A
Therefore, B

if they are going to cope with knight-and-knave puzzles. Unlike Rips, we believe that the ability to make these straightforward deductions depends not on formal rules of inference but on the ability to construct models of states of affairs. We have described this theory and a computer implementation of it elsewhere (see Johnson-Laird, Byrne, & Schaeken, 1989), and so we will sketch only its outlines here. The deduction above is made by constructing a set of models to represent the meaning of the first premise:

A	−B
−A	B
A	B

where each line represents a separate model. The information from the second premise can be incorporated only by eliminating those models that are inconsistent with it. The process leaves behind only one model:

−A	B

A procedure for formulating conclusions that do not correspond to any of the premises then yields the conclusion:

Therefore, B

For simplicity, we have used completely explicit models in this example. In fact, our theory assumes that ordinary individuals make explicit as little information as possible in their initial models of premises. They represent the disjunction above by the following models:

A
B

and they render these models wholly explicit only if forced to do so by the inferential task. The theory predicts that the greater the number of explicit models that have to be constructed, the harder the inferential task will be. This prediction has been confirmed in a number of experiments.

Granted a model theory of ordinary propositional deduction, how do people proceed with meta-logical puzzles? We believe that their ordinary deductive machinery does not cope with them. But, they possess a higher-level ability to reflect on these problems (and on their own processes of thought) using their ordinary deductive machinery as a sub-component. They typically have no existing procedures for dealing with meta-logical relations, and so their first efforts are tentative: they may, like a logician, pursue the consequences of certain assumptions about the truth or falsity of premises, they may notice certain interesting patterns in a puzzle, or they may grasp the consequences of circular assertions. With experience of the puzzles, they are likely to develop more systematic strategies – perhaps as a result of the "chunking" mechanism postulated by Newell (in press). Thus the development of meta-logical skill is a high-order inferential ability that depends on existing deductive procedures. We will now describe five meta-logical strategies, which all employ the model-based procedure for ordinary propositional deduction. And we will then show how the theory leads to certain predicted patterns of difficulty that are corroborated by the results of Rips's experiments.

Five strategies of meta-logical reasoning

We have modelled various meta-logical strategies by adding a component to our program that carries out model-based propositional deductions. The program parses a set of premises, such as:

A asserts that not A and not B
B asserts that B and A

and builds up a set of models representing their meaning. Because the meta-logical procedures use the ordinary deductive procedures as a sub-component, it is necessary to translate "true" and "false" into the language of that sub-component. Hence, the program uses the same symbol for both negation and falsity. This convenience, however, should not be taken to imply that we favour a general elimination of truth and falsity. In many cases, the translation is inadmissible, for example in the definition of validity.

In order to examine the plausibility of Rips's strategy, we have implemented a notational variant of it using models instead of formal rules of inference: the program follows up to the bitter end the *full chain* of consequences of making the contrasting pair of assumptions about an individual assertor. We will present two examples of its use, before we consider other less demanding strategies. The first problem (number 9 in the Appendix) has the premises:

A asserts that not B and C
B asserts that not A
C asserts that B

The chain from assuming the truth of A is straightforward, and yields a contradiction. It calls for the conjunction of the model representing that A is true:

A

with the models representing A's assertion:

−B C

The semantics of conjunction calls for forming the Cartesian product of the two sets of models, eliminating any inconsistencies:

A −B C

The program then follows up −B. The semantics of negating a proposition calls for constructing the complement of its set of models. The set representing B's assertion contains one model:

−A

and so its complement is obviously: A. This is consistent with the result above. Next, the program follows up C. The model of C's assertion is:

B

which yields a contradiction when it is conjoined with the result above. The output of the program, which summarizes the process thus far, is:

CHAIN hyp A → A −B C, neg-hyp B → −B A, hyp C → C B, ⇒ A CONTRADICTION

The chain from assuming the falsity of A is extremely complicated. The root of the problem is the need to follow up a set of disjunctive models. In assuming the falsity of A, the reasoner has to form the negation of A's assertion, that is, the negation of: not B and C. This negation, conjoined with not A, yields the following disjunctive set of models:

−A	B	C
−A	−B	−C
−A	B	−C

The procedure must now follow up each of the consequences within each of these models: hypothesizing B and then C in the first model, not B (which immediately yields a contradiction) in the second model, and B and then not C in the third model. The outcome is: −A B C. This model is passed to the procedure that describes models, and the result is the conclusion:

Not A and B and C

The program's output for this part of the chain is:

CHAIN neg-hyp A → −A B C, −A −B −C, −A B −C
DISJUNCTIVE CONSEQUENCES hyp B → B −A, hyp C → C B
DISJUNCTIVE CONSEQUENCES neg-hyp B → −B A, hyp B → B −A, neg-hyp C → −C −B ⇒ −A B C
IT FOLLOWS FROM THE CHAIN THAT NOT A AND B AND C

We will spare the reader the details of what happens when the procedure continues by constructing chains from the other two premises of the problem, though they also yield the same conclusion. In our view, the need to follow up the consequences of disjunctive models renders the strategy most improbable in the present case, especially granted that the problem was among the easiest for Rips's subjects (29% correct).

The second problem (number 19 in the Appendix) has the premises:

A asserts that not B
B asserts that A and C
C asserts that not A

The program yields the following output from the full chain based on the first premise:

CHAIN hyp A → A −B, neg-hyp B → −B −A C, −B A −C, −B −A −C, neg-hyp C → −C A ⇒ A −B −C
CHAIN neg-hyp A → −A B, hyp B → B A C ⇒ A CONTRADICTION
IT FOLLOWS FROM THE CHAIN THAT A AND NOT B AND NOT C

The use of the complete chain in this case is therefore less complicated than in the previous example. There is no need to follow up any disjunctive consequences. Yet, the problem was one of the most difficult that Rips investigated, and only 12% of the subjects solved it.

The contrast between these two examples led us to doubt whether Rips's strategy was correct (regardless of whether it was based on formal rules or mental models). We therefore developed four simpler strategies (based on informal protocols that we have collected).

A *simple chain* is constructed in the following way. Reasoners follow up the consequences of assuming the truth of the first premise, but they abandon the strategy whenever it becomes necessary to follow up disjunctive consequences. If they have not been forced to abandon the strategy, they then consider the consequences of assuming the falsity of the first premise, again abandoning the strategy whenever it is necessary to follow up disjunctive consequences. A further difference from Rips's full chain is that the procedure does not go on to consider the consequences of other premises.

A circular assertion, such as:

A asserts that A is false and B is true

catches the attention of most people. As our protocols show, reasoners are likely to grasp that the assertion appears to be self-refuting. Many people at this point can go no further. Some, however, grasp that if A's assertion is false, and A's assertion is a conjunction, then the first clause is true, and so the second clause must be false. This *circular strategy* accordingly first assumes that the assertor is telling the truth, and follows up only the immediate consequence of this assumption, i.e. it does not consider the consequences of this consequence (unlike the full chain). Next, it assumes that the assertor is making a false assertion, and follows up only the immediate consequence of this assumption. The circular strategy solves a problem if and only if one of these two assumptions leads to a contradiction and the other leads to an assignment of a truth value to all the individuals in the problem.

Here is an example. Given the problem:

A asserts that not A and not B
B asserts that B and A

the program produces the following output:

hyp A → NIL neg-hyp A → −A B ⇒ NOT A AND B

The output shows that the result of assuming the truth of A was nil, that is, a contradiction, and the result of assuming the falsity of A is a single model −A B. For problems that do not contain any circular assertions, this simple strategy is impotent.

So far, the three strategies that we have described all rely on making hypothetical assumptions and then following up their consequences to various degrees. There is an alternative tactic, however, that can be used once one has discovered that a particular hypothesis leads to a contradiction. Consider the problem (number 2 in the Appendix) that has the premises:

A asserts that A and B
B asserts that not A

There is a circular assertion, but the circular strategy fails, because it is necessary to trace two links (from A to B, and from B back to A, in order to discover that A must be false). Likewise, the simple chain fails because it is necessary to follow up the disjunctive consequences of negating A, i.e. −A B, −A −B. Once −A has been discovered as a consequence of hypothesizing A, there is a simple matching tactic that is a direct consequence of comparing models: −A is the case, and the content of B's assertion is −A, and so B must be true.

We have implemented this *hypothesize-and-match* strategy using again the model deductive procedure to carry out the essential work. The strategy considers the consequences of assuming that the first assertion, A, is true. If and only if the consequences lead to a contradiction, it then attempts to match −A to the content of the other assertions. If some other assertion, B, has a matching content, then B is true. This consequence can in turn be matched with the content of other assertions, and so on. It is possible to implement a mismatch tactic in which the falsity of B is derived from its inconsistency with some known truth, but, once again, we believe that such a strategy is likely to be beyond the competence of most people.

One other simple strategy also makes use of matching, and it is likely to be developed from encounters with the following sorts of premises (problem 5 in the Appendix):

A asserts that not C
B asserts that not C
C asserts that A and not B

Reasoners may notice that since A and B make the same assertion, they are either both true or both false. C, however, does not assign the same status to both of them. Hence, C is false. Both A's assertion and B's assertion match this conclusion, and so both are true. There are two tactics underlying this *same-assertion-and-match* strategy: first, the detection that two assertors make the same assertion, which in turn is inconsistent with a third assertion; and, second, the use of a match between the resulting conclusion (−C) and the content of specific assertions (A and B both assert not C). The strategy also detects where two individuals make opposing assertions about the same individual, and assigns falsity to any assertion that treats the two individuals as of the same status. We have little doubt that reasoners develop still other strategies depending upon the particular problems that they encounter, but we have not yet attempted to model them.

In summary, we have described five strategies:

1. Rips's full chain: assume that an assertor tells the truth, and follow up the consequences, and the consequences of the consequences, and so on. Assume that an assertor tells a lie, and do likewise. Carry out both processes for all premises.
2. Simple chain: assume that the assertor in the first premise tells the truth, and follow up the consequences, but abandon the procedure if it becomes necessary to follow up disjunctive consequences. Assume that the assertor in the first premise is lying and do likewise.
3. Circular: if a premise is circular, follow up the immediate consequences of assuming that it is true, and then follow up the immediate consequences of assuming that it is false.
4. Hypothesize-and-match: if the assumption that the first assertor A is telling the truth leads to a contradiction, then attempt to match −A with the content of other assertions, and so on.
5. Same-assertion-and-match: if two assertions make the same claim, and a third assertor, C, assigns the two assertors to different types, or vice versa, then attempt to match −C with the content of other assertions, and so on.

The predictions of the model theory

The four *simple* strategies that we described in the previous section are all based on the assumption that ordinary individuals have a limited ability to process models of premises. Hence, they cannot cope with negated conjunctions that force them to consider the consequences of a disjunctive set of models, they have only a limited ability to follow up the consequences of assumptions, and they find positive matches easier than negative mismatches. The model theory accordingly makes three main predictions about performance with meta-logical puzzles granted some minimal competence with them.

The first prediction is that problems that can be solved by using one of the simple strategies will be easier than those that require more powerful strategies such as the full chain proposed by Rips. In order to test this prediction, we have re-analysed the results of Rips's first experiment, which he kindly made available to us. The 34 problems are stated in the Appendix, and we have indicated those problems that can be solved using three of the simple strategies (simple chain, hypothesize-and-match, and same-assertion-and-match). We have ignored the circular strategy because the problems it solves are also soluble using a simple chain. Overall performance in the experiment was low (20% correct). Yet, there were 28% correct conclusions to the problems that can be solved by one of the four simple strategies, and only 14% correct conclusions to the problems that cannot be solved in this way (Mann-Whitney $U = 7$, $p < .001$, one-tail, by materials). This result was corroborated by separate analyses of the simple strategies versus more complex versions of them: simple chain problems 28% correct versus full chain problems 14% correct, $U = 2$, $p < .001$; hypothesize-and-match problems, 27% correct versus full matching problems, including those requiring negative mismatches, 12% correct, $U = 3.5$, $p < .002$. Too few problems could be solved by the similarity strategy to justify a statistical analysis.

The second prediction is that the difficulty of a problem will be a function of the number of clauses that it is necessary to use in order to solve the problem. This number obviously relates to the number of clauses in the statement of the problem, but the two notions are distinct, as we can illustrate by considering two contrasting examples.

The first problem has the premises:

A asserts that not A and B
B asserts that B

The circular strategy applied to the first premise yields the conclusion that A is false and hence B is false. The program in effect merely traverses the circular loop from A back to A in order to discover the contradiction. The consequences of B's assertion can be followed up, but they play no part in discovering the solution. The second problem has the premises:

A asserts that A and not B
B asserts that A

In this case the circular strategy fails to solve the problem, because it is necessary to consider both premises. The hypothesize-and-match strategy proceeds as follows: assume that A is true, and it follows that not B. From not B, it follows that not A, which contradicts the assumption. Given not A, B is not true because B asserts that A. Hence, the solution is: not A and not B. This problem should therefore be harder than the first one. The program in effect traverses the link from A to B, and then the link from B back to A, in order to discover the contradiction.

This second prediction, as we have illustrated, can be couched in terms of the number of links that have to be traversed in order to solve a problem. In this sense, the prediction is almost independent of the processing theory that we have proposed, and is likely to be made by any sensible analysis of meta-logical problems. Indeed, we suspect that the number of links to be traversed is one of the main contributors to the number of steps that Rips's program requires in order to solve a problem. The two problems above are indeed taken from his second experiment, and show how in this case our rather simpler account makes the same predictions as his theory. There is a corollary to our prediction. For many problems, the number of clauses (links) that have to be explored depends on the particular premise with which the process of reasoning begins. Hence, an experimental manipulation of this variable should affect performance.

The third prediction is a consequence of the model theory. Other things being equal, the hypothesis that an assertion is true should be easier to process than the hypothesis that an assertion is false: the operation of negating a set of models takes work. It calls, as we have seen, for the construction of the complement of a set of models. This prediction is corroborated by a result from Rips's second experiment, which by his own account presented some difficulty for his theory. The finding was that certain ways of couching a problem, such as:

A: I am a knave or B is a knight
B: I am a knight

are easier than others, such as:

A: I am a knave or B is a knave
B: I am a knight

The second version, however, calls for a slightly more complex process. Our simple chain strategy yields the following output for the first problem:

hyp A → A B, hyp B ⇒ B ⇒ A B
neg-hyp A → NIL ⇒ A CONTRADICTION

It yields the following output for the second problem:

hyp A → A −B, neg-hyp B → −B ⇒ A −B
neg-hyp A → NIL ⇒ A CONTRADICTION

As the reader will note, the second problem requires a negative operation that is not required by the first problem. This difference runs through the complete set of problems and accounts for the difference between them.

General discussion

Psychology is a "recursive" discipline because a plausible theory of high-level cognition should reveal how the theory itself could have been created as a result of high-level cognition. Hence, a theory of meta-logical deduction should provide some insight into its own development. Our theory postulates a capacity to think about the truth and falsity of premises (or indirect reflections of them in the guise of knights and knaves), which in turn depends on a general meta-cognitive capacity to reflect about problems and processes of thought at a higher level (see Johnson-Laird, 1983, Ch. 16). In this way, relatively simple reasoning strategies can be invented by logically untutored individuals. The same component can be used by logicians to create formal calculi for deduction, and then to reflect upon the relations between these calculi and their semantics. And, most importantly, it can be used by cognitive scientists to construct theories about itself. Given only Rips's single deterministic procedure, then neither these abilities nor the experimental results can be fully explained.

The major constraint that we have imposed on the meta-logical component is that the strategies that it develops are restricted by the capacity of working memory. Hence, the four simple strategies follow up the consequences of assumptions only to a limited extent, and they make positive matches more readily than negative mismatches. Yet they account for more aspects of Rips's first experiment than does his own theory. They also account for the two principal results of his second experiment. Indeed, they make additional more fine-grained predictions that we have not reported, because of lack of space. In short, we have presented a model theory of meta-logical reasoning that is explicit, that is not a notational variant of the natural deduction theory, and that accounts for the known phenomena of meta-logical deduction.

Envoi

Rips mounts an ingenious defence of his theory. He considers the following proposition:

If I am telling the truth, then the natural deduction theory is correct

As he points out, this assertion appears to be logically true (cf. Barwise and Etchemendy, 1987, p. 23). Alas, he here violates his own well-known principle, to which we have already alluded: "cognitive psychology has to do without semantic notions like truth" (Rips, 1986). He must therefore withdraw the argument on pain of contradiction. We recommend to him, however, the following robust Gödelian argument. Suppose that Rips ultimately devises a comprehensive theory of meta-logical reasoning based on formal rules. We will call this system, Rips Logic. Now consider the following proposition:

A asserts that A is not derivable in Rips Logic

This claim is either true or false. If it is false, then Rips Logic allows one to derive an assertion that asserts of itself that it is *not* derivable in Rips Logic. Hence, Rips Logic is inconsistent. If the claim is true, there is a true meta-logical assertion that cannot be captured within Rips Logic. Hence, Rips Logic is incomplete. Hence, a formal theory of meta-logical reasoning is bound to be either inconsistent or incomplete.

References

Austin, J.L. (1970). Truth. In J.O. Urmson & J.G. Warnock (Eds.), *Philosophical papers of J.L. Austin*, 2nd Edn. Oxford: Oxford University Press.

Barwise, J., & Etchemendy, J. (1987). *The liar: an essay in truth and circularity.* New York: Oxford University Press.

Boolos, G., & Jeffrey, R. (1980). *Computability and logic*, 2nd Edn. Cambridge: Cambridge University Press.

Johnson-Laird, P.N. (1983). *Mental models: Towards a cognitive science of language, inference, and consciousness.* Cambridge: Cambridge University Press.

Johnson-Laird, P.N., Byrne, R.M.J., & Schaeken, W. (1989). Reasoning by model: The case of propositional inference. Manuscript submitted for publication.

Newell, A. (in press). *Unified theories of cognition.* Cambridge, MA: Harvard University Press.

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

Rips, L.J. (1986). Mental muddles. In M. Brand and R.M. Harnish (Eds.), *The representation of knowledge and belief*, pp. 258–286. Tucson: University of Arizona Press.

Rips, L.J. (1989). The psychology of knights and knaves. *Cognition, 31*, 85–116.

Appendix. The results of Rips's Experiment 1, together with three simple strategies that can solve some of the problems

Note: We have stated the problems, not in the guise of knights and knaves, but in abbreviated form in which, for example, "A: −A & B" corresponds to A asserts that A is a knave and B is a knight. A cross in a column indicates that the corresponding problem can be solved by the corresponding strategy, where "same" refers to the same-assertion-and match-strategy, and "hypo-match" refers to the hypothesize-and-match strategy. All problems can be solved by Rips's full chain strategy provided that it is equipped with the appropriate rules of inference.

				Strategies		
Problem number	Problem	Solution	Percent correct	Same	Hypo-match	Simple chain
3 clauses						
1.	A: −A v −B					
	B: −A	A & −B	27			+
2.	A: A & B					
	B: −A	−A & B	27		+	
3.	A: −A & −B					
	B: −A	−A & B	27		+	+
4.	A: A v B					
	B: −A	A & −B	18			
4 clauses						
5.	A: −C					
	B: −C					
	C: A & −B	A & −C & B	35	+		
6.	A: B & −C					
	B: −A					
	C: B	−A & B & C	30		+	
7.	A: −A & − −B					
	B: A & B	−A & −B	30			+
8.	A: −A & −B					
	B: B & −A	−A & B	29			+
9.	A: −B & C					
	B: −A					
	C: B	−A & B & C	29		+	
10.	A: C					
	B: −C					
	C: A & B	−A & −C & B	27	+		
11.	A: −C					
	B: A v C					
	C: −B	−C & B & A	27			+

Problem number	Problem	Solution	Percent correct	Strategies: Same	Strategies: Hypo-match	Strategies: Simple chain
12.	A: C B: −C C: A & B	B & −C & −A	27	+		
13.	A: B & −C B: C C: −A	−A & B & C	24		+	
14.	A: B B: −C v A C: A	A & B & C	24			
15.	A: −B v −C B: −C C: A	A & −B & C	21			
16.	A: −B & C B: −A C: B	−A & B & C	21		+	
17.	A: B & C B: −C C: A	−A & B & −C	15			
18.	A: −B & C B: A C: B	−A & −B & −C	12			
19.	A: −B B: A & C C: −A	A & −B & −C	12			
20.	A: A B: A C: A & −B	−C (A ↔ B)	0			
5 clauses						
21.	A: B B: C & A C: −A & B	−A & −B & −C	27			
22.	A: −B B: A v C C: −A & C	−A & B & C	27			+
23.	A: A v − A B: −C & −A C: −B	A & −B & C	21			
24.	A: A v −A B: −C & −A C: B	A & −B & −C	18			
25.	A: B v −B B: −C C: −A & −B	A & B & −C	18			

				Strategies		
Problem number	Problem	Solution	Percent correct	Same	Hypo-match	Simple chain
26.	A: C & A B: A & −C C: −A	−A & −B & C	15			
27.	A: B & A B: B & −C C: A	−A & −C (B v −B)	9			
28.	A: −B & −C B: −A & −C C: − −B	A & −B & −C	6			
29.	A: −B & A B: A & C C: −B	−A & −B & C	6			
30.	A: B & C B: B & −A C: A	−A & −C (B v −B)	3			
9 clauses						
31.	A: −A & −B & −C B: A v B v C C: −A v −B v −C	−A & B & C	29		+	
32.	A: A & B & C B: −A v −B v −C C: A v B v C	−A & B & C	27		+	
33.	A: −B & −C & −A B: −C & −B & A C: −A & −B & C	−A & −B & C	21			
34.	A: B & C & A B: C & B & −A C: −A & −B & C	−A & −B (C v −C)	0			