

The Psychology of Syllogisms

PHILIP N. JOHNSON-LAIRD

Sussex University

AND

MARK STEEDMAN

Warwick University

Two experiments were carried out in which subjects had to draw conclusions from syllogistic premises. The nature of their responses showed that the figure of the syllogisms exerted a strong effect on the accuracy of performance and on the nature of the conclusions that were drawn. For example, premises such as "Some of the parents are scientists; All of the scientists are drivers" tend to elicit the conclusion, "Some of the parents are drivers" rather than its equally valid converse, "Some of the drivers are parents". In general, premises of the form $\frac{A-B}{B-C}$ created a bias towards conclusions of the form $A-C$, whereas premises of the form $\frac{B-A}{C-B}$ created a bias towards conclusions of the form $C-A$. The data cast doubt on current theories of syllogistic inference; a new theory was accordingly developed and implemented as a computer program. The theory postulates that quantified assertions receive an analogical mental representation which captures their logical properties structurally. A simple heuristic generates putative conclusions from the combined representations of premises, and such conclusions are put to logical tests which, if exhaustively conducted, invariably yield a correct response. Erroneous responses may occur if there is a failure to test exhaustively.

Only connect.

—E. M. Forster

The first experimental investigation into syllogisms appears to have been carried out about 70 years ago by Störring (see Woodworth, 1938), and since then there has been a steady series of studies of the various factors

We are grateful to Janellen Huttenlocher for her very considerable advice, and to Paul Byerly, Diana Shapiro, and David Weltman for help in conducting the experiments. We also thank the editor, the anonymous referees, and Stuart Sutherland for their helpful comments on expository matters. The research was carried out, in part, with grants for scientific assistance from the Medical Research Council and the Social Science Research Council (GB). Reprint requests should be sent to P. N. Johnson-Laird, Laboratory of Experimental Psychology, University of Sussex, Brighton BN1 9QG, England.

affecting their difficulty. Yet, it is only in the last few years that any model of the complete inferential process has been proposed. Part of the reason for such slow progress would seem to be the baleful influence, Aristotle excluded, of traditional logic. The scholastics recognized 64 different moods of syllogism, since the two premises and the conclusion have to be selected from the four moods of sentence ($4 \times 4 \times 4 = 64$):

- | | |
|----------------------------|-------------------|
| (A) Universal affirmative | All A are B |
| (I) Particular affirmative | Some A are B |
| (E) Universal negative | No A are B |
| (O) Particular negative | Some A are not B, |

where the parenthesized letters are the traditional mnemonics (derived from *Affirmo* and *Nego*). Traditional logic also admits four different "figures":

First figure	Second figure	Third figure	Fourth figure
M—P	P—M	M—P	P—M
S—M	S—M	M—S	M—S
$\therefore$ S—P	$\therefore$ S—P	$\therefore$ S—P	$\therefore$ S—P ,

where S denotes the subject of the conclusion, P denotes the predicate of the conclusion, and M denotes the middle term which is common to both premises but disappears in the deduction.

If there are 64 moods and four figures, it is natural to suppose that there are $64 \times 4 = 256$ different sorts of syllogism. This number has certainly been taken for granted by many psychologists. It is, of course, erroneous: There are twice that number of syllogisms. Logicians ignored the order of the premises and made an arbitrary decision to cast their figures so that the subject (S) of the conclusion always occurs in the second premise. Obviously, logic is not affected if S happens to occur in the first premise, e.g.,

$$\begin{array}{c}
 S-M \\
 M-P \\
 \hline
 S-P
 \end{array}$$

and it is worth noting that Aristotle used a set of figures that included this "deviant" case. One immediate consequence of the slavish adherence to scholastic logic is a general neglect of half the possible syllogisms, and this omission has had serious consequences for understanding the psychology of syllogisms.

Another reason for slow progress has to do with experimental technique. Rather than an attempt to castigate the shortcomings of particular

investigations, it will be simpler to spell out some of the desirable requirements of an experiment on syllogistic reasoning:

First, subjects should have to make a deduction in order to carry out their experimental task. The point is obvious, but the drawbacks of conventional techniques are more subtle. If the task is merely to evaluate a given syllogism as valid or invalid, a subject may carry it out without ever having to make an inference. Even a multiple choice between different putative conclusions may tend to obscure the deductive process, either because of the particular set of alternatives chosen by the experimenter or because of some idiosyncratic procedure that a subject adopts, such as working backward from conclusion to premises, guessing the most plausible conclusion, and so on.

Second, subjects should be given a representative selection of problems. It is little use seeking to draw general conclusions on the basis of, say, a dozen syllogisms when the total possible number is 512.

Third, syllogisms should be presented with a sensible, though non-controversial, linguistic content. Although a case has been made in the past for studying inference with an abstract or symbolic content, it is now known that such materials can lead to qualitative changes in performance (Wason and Johnson-Laird, 1972). A psychologist who studies reasoning with abstract materials is not so much studying a pure deduction, unsullied by his subjects' knowledge or attitudes, as a very special sort of reasoning designed to compensate for the absence of everyday content.

Fourth, in the analysis and description of results, it is crucial to consider each syllogism separately. A number of published studies present only data pooled across different figures or across different moods. Such an exposition may be appropriate for the evaluation of an author's own hypothesis, but it can render the data useless for anyone who wishes to examine an alternative theory or to construct a general model of syllogistic inference.

These four simple requirements have never been satisfied by any investigation to be found in the literature, and even some recent studies have neglected them. Our initial goals were accordingly to carry out an experiment in which subjects drew their own inferences from a reasonable selection of sensible syllogisms and to try to give an account of how they performed this task.

EXPERIMENT 1: VALID SYLLOGISMS

The aim of the experiment was to examine our subjects' ability to make valid syllogistic deductions. The subjects were presented with pairs of syllogistic premises and asked to state in their own words what conclusion followed logically from them. This technique has a subsidiary advantage: Although there are 512 different syllogisms, there are only 64 different pairs of premises, and this reduction makes it very much

more feasible to test a representative sample of problems. In the present experiment, subjects were tested with the 27 pairs of premises that yield a valid conclusion; that is, at least one of the eight possible conclusions is correct.

Design and Materials

Each subject was asked to make a deduction from the 27 pairs of premises that are shown in Table 1 with their valid conclusions italicized. The problems were presented with a sensible content of a sort unlikely to predispose subjects toward a particular conclusion. Hence, a typical pair was:

None of the musicians are inventors
All of the inventors are professors
∴ ?

The materials were mimeographed and assembled into booklets in different random orders.

Procedure

The subjects were tested individually. There were told that they were going to take part in an investigation of the way in which people combine information in order to draw conclusions from it. They would be given a series of pairs of statements about people whom they were to imagine as assembled in a room. Their task was to write down what followed from each pair of statements about the occupants of the room. The purpose of this instruction was to insulate still further the content of the problems from subjects' attitudes or expectations. The subjects were also instructed that their answers were to be based solely upon what could be deduced with absolute certainty from the premises, and it was made clear that for every problem there was always at least one such conclusion that could be drawn. The subjects were allowed as much time as they wanted in order to complete the task.

Subjects

Twenty volunteers, who were undergraduate students at University College London, participated in the experiment.

Results

The best way to illustrate the conventions that we have adopted in presenting the results is to examine in detail the actual performance on one problem. Consider the following premise pair and the 14 valid responses it evoked (we present premises and conclusions in an abstract form for ease of reading):

All of the A are B	
All of the B are C	
∴ <i>All the A are C</i>	(seven subjects)
∴ <i>All A are C</i>	(five subjects)
∴ <i>The A are C</i>	(one subject)
∴ <i>All A are B and C</i>	(one subject)

These conclusions despite some superficial heterogeneity are all logically impeccable; here, and throughout the paper, we italicize valid deductions. (The last response is interesting because it is an instance of a partially digested middle term.) Since we shall be concerned only with logic and the order of terms, we shall ignore superficial variations and classify all of these responses under the same general rubric: *All A are C*. No claims about the present experiment hinge upon any more subtle distinctions in the data.

Table 1 presents the frequencies of the main responses (i.e., any response made by two or more subjects) to each of the 27 problems. Since 20 subjects participated, the residual untabulated value for any problem corresponds to those miscellaneous erroneous responses that were not made by more than one subject.

What is immediately evident from a casual inspection of Table 1 is the variation in the difficulty of the problems. With the easiest premise pair 17 out of 20 subjects produced valid deductions; with the hardest premise pair only six out of 20 subjects produced valid deductions. This aspect of the results was reinforced by the extent to which subjects agonized over conclusions: Sometimes a conclusion would emerge rapidly, sometimes only a very hesitant conclusion emerged which was often either erroneous or merely a restatement of the premises. The variation is remarkable but well established, and we will make no further comment on it at this point. There is a more important phenomenon to be considered.

When the results in each quadrant of Table 1 are examined, it is evident that there is a pronounced "figural" effect. Thus, the following examples show a highly reliable bias toward one of two equally valid conclusions:

Some A are B All B are C	All B are A Some C are B
$\therefore$ <i>Some A are C</i> (15 subjects)	$\therefore$ <i>Some C are A</i> (16 subjects)
$\therefore$ <i>Some C are A</i> (2 subjects)	$\therefore$ <i>Some A are C</i> (1 subject).

Such patterns are entirely representative: With the $\frac{A}{B} - \frac{B}{C}$ figure (the top left-hand quadrant), 71% of all the valid conclusions were of the form $A - C$, whereas with the $\frac{B}{C} - \frac{A}{B}$ figure (the bottom right-hand quadrant), 70% of the valid conclusions were of the form $C - A$. The bias is evident both in those syllogisms that permit two converse conclusions and in the ease of solution to those premises that permit only a single conclusion and not its converse. There was little if any bias in the case of the $\frac{A}{C} - \frac{B}{B}$ figure (53% valid conclusions of the form $C - A$) and no bias at all in the case of the $\frac{B}{B} - \frac{A}{C}$ figure. All 20 subjects had a bias toward $A - C$ conclusions for the

TABLE I
THE FREQUENCIES OF THE MAIN SORTS OF DEDUCTION IN EXPERIMENT 1^a

Second premise	First premise							
	All A are B	Some A are B	No A are B	Some A are not B	All B are A	Some B are A	No B are A	Some B are not A
All B are C	All A are C All C are A	14 2	Some C are not A No C are A Some A are C	15 2	Some C are not A No C are A Some A are C	7 3 3 All B are A and C 5	10 4 4 No A are C	Some C are not A Some C are not A Some C are not A Some C are not A
Some B are C			Some C are not A No A are C No C are A	9 3 2	Some A are C Some C are A No C are A	8 4	Some C are not A Some C are not A	
No B are C	No A are C No C are A Some A are C	13 2 2	Some A are not C Some A are C Some A are C	10 4	Some A are not C No C are A No A are C	8 3 2	Some A are not C Some A are C Some C are not A	
Some B are not C					Some A are not C Some A are not C Some A are not C	12 3	Some A are not C Some A are not C Some A are not C	
All C are B			No A are C No C are A No A are C	9 6	All C are A All A are C	12 4	No C are A No A are C	10 5
Some C are B			Some C are not A Some C are A	12 2	Some C are A Some A are C	16 1	Some C are not A Some C are A	12 3
No C are B	No C are A No A are C	11 6	Some A are not C Some C are A	11 2	Some A are not C No A are C No C are A	6 5 3	Some A are not C No C are A Some C are A	9 2 2
Some C are not B	Some C are not A Some A are not C Some C are B (sic)	12 2 2						

^a Valid deductions are italicized. The eight columns represent the first premise in a problem, the eight rows the second premise. The maximum total for any problem is 20: Where totals are less than 20, the residuals correspond to miscellaneous errors.

$\frac{A-B}{B-C}$ premises and a bias toward $C-A$ conclusions for the $\frac{B-A}{C-B}$ premises, and hence the effect is highly significant. The effect occurs for all the relevant syllogisms except those that have only one valid conclusion which is inconsistent with the bias.

Discussion

Why has such a striking phenomenon as the figural effect not been reported before? The answer is that the crippling effect of traditional logic has led to a neglect of syllogisms in the "unorthodox" figures; the stigma attached to such figures is, indeed, still observable in recent publications. The effect was, however, partially anticipated by Frase (1968) who observed that orthodox syllogisms in the first figure,

$$\begin{array}{c} B-A \\ C-B \\ \hline \therefore C-A, \end{array}$$

were more often evaluated correctly than orthodox syllogisms in the fourth figure,

$$\begin{array}{c} A-B \\ B-C \\ \hline \therefore C-A, \end{array}$$

with the other two figures yielding an intermediate performance. Doubtless, if Frase had used syllogisms with conclusions of the form $A-C$, the difference would have reversed and he would have established the complete figural effect. Frase explained his results by analogy with the mediational paradigms of paired-associate learning, since the first figure corresponds to a "forward chain" and the fourth figure corresponds to a "backward chain", but the analogy seems to break down in the light of the full figural effect. Both Wilkins (1928) and Sells (1936) used some syllogisms in unorthodox figures, and with hindsight one can also detect some traces of a figural effect in their data.

EXPERIMENT 2: VALID AND INVALID SYLLOGISMS

Although a figural effect emerged clearly from Experiment 1, we were unable to formulate a theoretical explanation of it in the absence of data from premises from which no valid deduction could be drawn. Hence, a second experiment was devised in order to try to replicate the results of the first one and to extend the technique to premises lacking valid conclusions. The experiment also allowed us to evaluate the main hypotheses about syllogistic inference, since they largely concern invalid inferences.

Design

The aim of the experiment was to assess performance with the complete set of syllogisms, and accordingly each subject attempted to make an inference from all 64 possible pairs of premises illustrated in Table 1. This task was performed twice by every subject on two separate occasions approximately a week apart. The contents of the syllogisms were similar to those used in Experiment 1, except that a more stringent attempt was made to minimize semantic relations between the terms within each premise pair while retaining moderate plausibility for them and for any conclusion, valid or invalid. We found that the most successful way to construct premises within these constraints was to use one term denoting an occupation and two terms denoting preoccupations or interests, for example, "All of the gourmets are storekeepers. All of the storekeepers are bowlers." Two separate lists of the 64 problems were constructed. In order to create the second list, the contents of the 27 soluble problems in the first list were exchanged with those of 27 insoluble problems, and the remaining 10 insoluble problems had their contents reassigned from one problem to another at random. The subjects received one list in the first session and the other list in the second session in a counterbalanced pattern. Each list was presented in a random order.

Procedure

The procedure was the same as in Experiment 1 except that the subjects' performance was timed and they were told to make their responses both accurately and as quickly as possible. They were also told to restrict their answers to one of the four moods or else to state that no valid conclusion followed from the premises.

Subjects

Twenty paid volunteers, students at Teachers College, Columbia University, were tested individually in the experiment.

Results

The results for each of the 64 problems are presented in the four tables in the Appendix. In order to simplify the presentation we shall mainly consider performance in the second test. However, the pattern of the results was with one exception (to be discussed below) very similar in both tests, as the reader may care to check, and the significant effects we report are manifest in the results from the first test.

The Effect of Figure on the Form of Conclusions

By far the most important result is the confirmation of the figural effect. In the case of the $\begin{smallmatrix} A \\ B \end{smallmatrix} \text{---} \begin{smallmatrix} B \\ C \end{smallmatrix}$ figure there was a strong bias towards a conclusion of the form $A \text{---} C$, and in the case of the $\begin{smallmatrix} B \\ C \end{smallmatrix} \text{---} \begin{smallmatrix} A \\ B \end{smallmatrix}$ figure, there was a strong bias towards a conclusion of the form $C \text{---} A$. (The phenomenon has turned out to be easy to replicate: When audiences at universities as far afield as Chicago, New York, Edinburgh, London, Padova, and Nijmegen, were presented with appropriate syllogisms, they all showed a massive figural effect in their conclusions.) The present data are summarized in Table 2, but the results of the second test will be analyzed in more detail.

TABLE 2

THE OVERALL EFFECT OF FIGURE ON THE FORM OF CONCLUSIONS IN EXPERIMENT 2:
THE PERCENTAGES OF A—C AND C—A CONCLUSIONS AS A
FUNCTION OF THE FIGURE OF THE PREMISES

Form of conclusion	Figure of premises			
	A—B	B—A	A—B	B—A
	B—C	C—B	C—B	B—C
A—C	(51.2) ^a 44.7	(4.7) 5.3	(21.2) 13.7	(31.9) 29.4
C—A	(6.2) 7.8	(48.1) 45.3	(20.6) 28.1	(17.8) 25.0

^a Numbers in parentheses are from the first test.

When a syllogism has only one valid conclusion (i.e., the converse conclusion is invalid), then subjects tend to make this deduction where it is compatible with the figure of the premises, but tend to be unable to make it where it is incompatible with the figure of the premises. Thus, the following two examples of IE and EI problems illustrate the relative ease of drawing a conclusion compatible with the figure

Some A are B No B are C	No B are A Some C are B
∴ Some A are not C (17 subjects)	∴ Some C are not A (14 subjects),

whereas the following two examples illustrate the relative difficulty of drawing a conclusion incompatible with the figure

No A are B Some B are C	Some B are A No C are B
∴ Some C are not A (8 subjects)	∴ Some A are not C (5 subjects).

The overall differences are massive: Where valid deductions were compatible with figure, they occurred on 85% of occasions with the $\frac{A-B}{B-C}$ figure and 77.5% of occasions with the $\frac{B-A}{C-B}$ figure, but where they were incompatible with the figure they occurred on only 20% of occasions. There was not a single exception in all 20 subjects to this pattern of results. There was a comparable bias in the invalid conclusions that occurred for these problems, though the frequencies were too low for statistical comparison.

There was also a pronounced figural effect for syllogisms with two

(converse) valid conclusions: 82.5% conclusions compatible with figure versus 10% incompatible with it for $\frac{A-B}{B-C}$ premises, and 77.5 versus 7.5% for $\frac{B-A}{C-B}$ premises. (The effect is present in the results of 18 subjects, absent in the results of one subject, and controverted by the results of one subject; Sign test, $p < 0.0005$, two tailed.) It is noteworthy that in none of these syllogisms was there ever an invalid conclusion incompatible with figure.

Finally, in the case of syllogisms lacking a valid conclusion, there was a figural effect in the invalid conclusions that were drawn: 30% invalid conclusions were drawn compatible with the $\frac{A-B}{B-C}$ figure, but only 2.5% were drawn incompatible with it, and 31.5% invalid conclusions were drawn compatible with the $\frac{B-A}{C-B}$ figure but only 3.5% were drawn incompatible with it. (This trend was also confirmed by the results of 18 subjects, disconfirmed by the results of one subject, and there was a tie for one subject.)

Turning to the other two remaining figures, there was some bias towards conclusions of the form $C-A$ for the $\frac{A-B}{C-B}$ figure. The bias was evident in the results for premises with only one valid conclusion (67.5% correct where the required conclusion was of the form $C-A$, but only 30% correct where it was of the form $A-C$; Wilcoxon test, $p < 0.01$, two-tailed), but the bias for premises with two (converse) valid conclusions was not statistically significant (45% of the form $C-A$ but 30% of the form $A-C$), and it was only very slight for premises with no valid conclusion (14% invalid conclusions of the form $C-A$, and 9% invalid conclusions of the form $A-C$). There did not appear to be any reliable bias towards one form of conclusion or the other in the case of the $\frac{B-A}{B-C}$ figure.

The Effect of Figure on Accuracy

There was a considerable difference in the difficulty of the problems (Cochran's $Q = 368$, with $df = 63$, $p < 0.001$). The difference reflects in part the figure of the premises ($\chi^2 = 10.8$, $df = 3$, $p < 0.025$, Friedman two-way analysis of variance). Table 3 states the percentages of correct responses for the four figures both for premises with a valid conclusion and for premises without a valid conclusion.

For each $\frac{A-B}{B-C}$ problem, there is a corresponding $\frac{B-A}{C-B}$ problem: In effect, the order of the two premises is simply reversed. The $\frac{A-B}{B-C}$ figure yielded a superior performance for all six of the problems with valid conclusions, and the difference is statistically significant. It was

TABLE 3
THE PERCENTAGES OF CORRECT RESPONSES FOR THE FOUR FIGURES

	Figure of premises				Overall
	A—B B—C	B—A C—B	A—B C—B	B—A B—C	
Premises with valid conclusions	(60) ^a 68	(50) 58	(53) 58	(49) 69	(53) 64
Premises with no valid conclusions	(52) 66	(53) 64	(71) 75	(74) 84	(61) 71
Overall	(55) 67	(52) 62	(64) 68	(60) 76	(58) 68

^a Numbers in parentheses are from the first test.

evident in the data of 11 subjects, four subjects yielded data in the opposite direction, and there was no difference between the figures for the remaining five subjects (Wilcoxon test, $p < 0.05$, two-tailed). However, there is no real difference between these two figures for problems lacking a valid conclusion. Comparisons with (and between) the other two figures can only be made globally because they do not contain equivalent problems. The impression that these other two figures show a marked superiority for premises lacking a valid conclusion is borne out in the data for all but three subjects, a highly significant difference (Sign test, $p < 0.003$, two-tailed). In fact, there is an interaction here: These two figures show a greater advantage for premises lacking valid conclusions (in comparison with those having a valid conclusion) than do the first two figures (14 subjects conform to this trend, six subjects contravene it; Wilcoxon test, $p < 0.05$, two-tailed). An interesting related result concerns the proportion of times subjects respond "No valid conclusion" to premises that, in fact, permit a valid conclusion to be drawn: Table 4 presents these data for the four figures. The figure of the premises here exerted a significant effect on the

TABLE 4
THE PERCENTAGES OF "NO VALID CONCLUSION" RESPONSES TO THOSE PREMISES WITH VALID CONCLUSIONS IN THE FOUR FIGURES^a

Figure of premises			
A—B B—C (<i>n</i> = 6)	B—A C—B (<i>n</i> = 6)	A—B C—B (<i>n</i> = 6)	B—A B—C (<i>n</i> = 9)
(13) 7	(21) 18	(22) 22	(23) 13

^a Data in parentheses are from the first test; *n* = number of premise pairs.

propensity to respond "No valid conclusion" (Friedman two-way analysis of variance, $\chi_r^2 = 9.4$, $df = 3$, $p < 0.05$).

The Effect of Mood

Since mood has long been known to have effects on syllogistic inference, our analysis of it in the present experiment will be brief. Table 5 presents the percentages of correct responses (from the second test) pooled in terms of the mood of the premises. It is evident from the degree of variation that mood has a marked effect on performance ($\chi_r^2 = 71$, $df = 9$, $p < 0.001$, Friedman two-way analysis of variance). The critical point, of course, is that the difficulty of a syllogism depends on both its figure and mood. It is evident from our discussion of the IE and EI examples in the first section of results that neither figure nor mood alone is sufficient to predict difficulty. The two variables interact. Their interaction not only determines the difficulty of a problem but also the characteristic conclusions that it elicits.

The Effect of Content and Practice

In order to simplify the analysis of the results, we have assumed that performance on the problems was independent and that the content of the problems had no systematic effect upon performance. These assumptions are in general supported by the results, for example, the consistency of the figural effect. However, there was a distinct improvement in performance from the first test to the second test. There were 58% correct responses in the first test and 68% correct responses in the second test. Nineteen out of the 20 subjects improved performance, and an improvement was registered on 44 out of the 64 premise pairs. One striking differential effect of practice concerned those pairs of premises in the $\frac{B-A}{B-C}$ figure that permit a valid conclusion to be drawn: There was an improvement from 49% correct on the first test to 69% correct on the

TABLE 5
THE PERCENTAGES OF CORRECT RESPONSES ON THE SECOND TEST
IN TERMS OF THE MOOD OF THE PREMISES

Mood of the second premise	Mood of the first premise			
	A	I	E	O
A	69	68	51	39
I	76	83	60	81
E	50	54	71	74
O	54	89	78	94

second test, with only two subjects failing to enhance their performance on these nine problems (Sign test, $p < 0.001$, two-tailed). Much of this improvement is due to a decline in the propensity to respond "No valid conclusion" (see Table 4).

Because the responses themselves proved to be extremely revealing as well as varying considerably in their accuracy, we shall report no extensive analysis of the latency data. It is difficult to treat the latencies statistically because of the diversity of responses to many pairs of premises. Moreover, there was a very reliable correlation between latency and accuracy: Those premises that yielded many correct answers also yielded them rapidly ($\tau = 0.37$, $p < 0.0001$, two-tailed).

THE ANALOGICAL THEORY OF REASONING WITH QUANTIFIERS

Because the present results cast considerable doubt on current theories, our major task is to provide a satisfactory explanation of syllogistic reasoning. This explanation must account both for invalid and valid deductions because competent adults make mistakes but are capable of rational thought under optimal circumstances. Indeed, without this ability, it is difficult to see how the logic of syllogisms could have been formulated in the first place.

The central assumption of the present theory (of which a preliminary account may be found in Johnson-Laird, 1975) is that syllogistic inference is based on an *analogical* representation of the premises that captures their logical properties within its structure. The theory postulates four stages in the process of inference:


- (1) a semantic interpretation of the premises,
- (2) an initial *heuristic* combination of the representations of the two premises,
- (3) the formulation of a conclusion corresponding to the combined representation, and
- (4) a *logical* test of the initial representation which may lead to the conclusion being modified or abandoned.

As a heuristic for generating potential conclusions, the theory postulates a bias toward forming connections during the process of combining the representations of the premises. The process of logical testing involves an attempt to break links that may have been invalidly forged in this way. We now consider the four stages in more detail.

Stage 1: The Interpretation of the Premises


One clue to the mental representation of quantified assertions was provided by a subject in Experiment 1. When he was asked to describe how he had performed the task, he replied, referring to a specific

premise, "I thought of all the little (sic) artists in the room and imagined that they all had beekeeper's hats on." This remark provided the germ of an idea for a new hypothesis about the semantic representation of quantified assertions: *A class is represented simply by thinking of an arbitrary number of its exemplars*. Thus, a subject represents a statement such as "All the artists are beekeepers," first by imagining an arbitrary number of artists, which he takes to represent a relevant class of them, and then by tagging each of them in some way as a beekeeper. Since there may be beekeepers who are not artists, he adds an arbitrary number of such beekeepers to his representation tagging them in some way as optional. The various elements in the representation may be vivid images or abstract or verbal items. What is important is not their phenomenal content but the relations between them. Accordingly, the representation of "All the artists are beekeepers" might have the following form:



This is a direct analog of the logic of the assertion: There are an arbitrary number of artists tagged as beekeepers, and the parenthetical items represent the possibility of an arbitrary number of beekeepers who are not artists. The arrows stand for the semantic relation of class membership (each artist *is a* beekeeper); they are directional because we assume that it is relatively easy to traverse the link from artist to beekeeper but relatively difficult, though not impossible, to traverse it in the opposite direction. What we have in mind (and indeed actually exploit in a computer implementation of the theory) is a list-structure. The representation of each artist has stored with it the address of the corresponding representation of the beekeeper, and it is consequently easy to move from artist to beekeeper, but the representation of a beekeeper has no concomitant address of an artist, and the only way to move from beekeeper to artist is to search through all the artists until an appropriate link is found that leads back to the starting place.

Although arbitrary numbers of exemplars are involved, for convenience we will illustrate the various sorts of premise with a minimal number. An A premise, "All A are B," has the representation:



An I premise, "Some A are B," has the representation



which allows for the possibility that there are a 's which are not b 's, and b 's which are not a 's.

An E premise, "No A are B," has the representation

$$\begin{array}{cc} a & a \\ \perp & \perp \\ b & b \end{array}$$

where the stopped arrows indicate negative links. It is, of course, insufficient to represent negation merely by the absence of a link because subsequent processes might lead to a link being established: The broken link, $a \text{ ---} \perp b$ corresponds to a negative which prohibits a positive link between the a and any b .

An O premise, "Some A are not B," has the representation

$$\begin{array}{cc} a & (a) \\ \perp & \downarrow \\ b & b \end{array} .$$


This allows for the possibility that there are a 's which are b 's, and for the possibility that no a 's are b 's; that is, when the optional a is omitted, the linked b drops out with it. There is evidence that "Some A are not B" is often taken to imply "Some A are B" (e.g., Johnson-Laird, 1970) though our representation leaves it as an option. In accordance with traditional terminology, any term that has an optional element in its representation is "undistributed," otherwise it is "distributed;" for example, the term A in "Some A are not B" is undistributed, whereas the term B is distributed.

Stage 2: The Heuristic Combination of the Representations of Premises

Some heuristic is required in order to generate putative conclusions because logic cannot determine what conclusions to draw but at most whether a *given* conclusion is valid. There is evidence from other tasks for a bias toward making one-to-one matches and toward verification (see, e.g., Wason and Johnson-Laird, 1972, p. 241). Hence, we assume that in combining the premises there is a heuristic bias toward forming thoroughgoing connections between all the classes, that is, a bias toward linking up end items by way of middle items. Such premises as

All the artists are beekeepers
Some of the beekeepers are chemists

will be combined in such a way that the beekeepers who are chemists will be chosen from among those who are artists. Thus, the combined representation has the following sort of structure:




This representation readily leads to the invalid conclusion: "Some of the artists are chemists" (a conclusion that 12 subjects actually drew). With negative premises, we assume that there is the same bias toward trying to link up end items by way of the middle items, but in such a case the path will be a negative one. In general, where a path contains two positive links, it is positive; where it contains at least one negative link, it is negative; and any other path is indeterminate; for example, $a \rightarrow b$ (c), because the missing link could be positive or negative. The theory makes no assumptions about the order in which paths are constructed or assessed.

Stage 3: The Formulation of a Conclusion

In order to formulate a conclusion, it is necessary to determine the nature of the paths between the end items in a representation. The logic of this process is transparent. Where there is at least one negative path, then the conclusion is of the form *Some X are not Y*, unless there are only negative paths in which case it is of the form *No X are Y*. Otherwise, where there is at least one positive path, the conclusion is of the form *Some X are Y*, unless there are only positive paths in which case it is of the form *All X are Y*. In any other case, no valid conclusion can be drawn, that is, where there are only indeterminate paths.

Let us consider some specific examples of the initial combinations of premises and the conclusions that would be drawn from them. We have indicated the appropriate interpretation of each path in the following examples, using "+" for a positive path, "-" for a negative path, and "?" for an indeterminate path. The first example is a simple valid deduction, with one positive path and one indeterminate path:



A strong figural effect is predicted because the paths are in a uniform direction; it is confirmed, and only three subjects produced the equally valid but nonoptimal conclusion. The next examples illustrate what happens when there is no optimal direction in which to establish a connection between the end items and consequently no predicted bias in the form of the conclusions:

	+	?	
All B are A	<i>a</i>	<i>a</i>	(<i>a</i>)
	↑	↑	∴ <i>Some A are C</i> (11 subjects)
Some B are C	<i>b</i>	(<i>b</i>)	
	↓		∴ <i>Some C are A</i> (9 subjects)
	<i>c</i>	(<i>c</i>)	
	+	?	
	—	—	—
No A are B	<i>a</i>	<i>a</i>	<i>a</i>
	⊥	⊥	⊥ ∴ <i>No A are C</i> (9 subjects)
All C are B	<i>b</i>	<i>b</i>	(<i>b</i>)
	↑	↑	∴ <i>No C are A</i> (6 subjects)
	<i>c</i>	<i>c</i>	
	—	—	

Stage 4: The Logical Test of an Initial Representation

An initial representation is formed on the basis of a heuristic. Once it is formed, however, it is possible to bring logic to bear on it. If the heuristic is analogous to a bias toward verification, then the logical test is analogous to an attempt at falsification: It consists in trying to break the established paths between end items without doing violence to the meaning of the original premises. Thus, in the case of the example, "All the artists are beekeepers; some of the beekeepers are chemists," the initial representation tags one of the beekeepers who is an artist as a chemist. In testing this representation, one can establish the link from a beekeeper that is *not* an artist to the chemists, and in this way destroy the path leading from artists to chemists. It follows that no conclusion can be read off from the modified representation. Hence, the initial representation


	+	?
All A are B	<i>a</i>	<i>a</i>
	↓	↓
Some B are C	<i>b</i>	<i>b</i> (<i>b</i>)
	↓	
	<i>c</i>	(<i>c</i>)
	+	?

is modified as the result of the logical test to


	?	?
All A are B	<i>a</i>	<i>a</i>
	↓	↓
Some B are C	<i>b</i>	<i>b</i> (<i>b</i>)
		↓
	<i>c</i>	(<i>c</i>)
	?	?

from which no valid conclusion can be derived.


The principles governing the testing of negative syllogisms reflect the same basic idea. The initial bias is toward establishing negative links, and testing consists in trying to break them or to establish a positive link where a negative one had prevailed, without doing violence to the meaning of the premises or creating a contradiction. Here is a typical example where an initial representation gives rise to a fallacious inference:




But such a representation readily permits the establishment of positive links incompatible with the initial conclusions:



A truly rational subject may have to pursue a vigorous search in order to establish the correct conclusion to a pair of premises. Consider, for instance, what could happen with premises of the form: "Some A are B," "No C are B". Their initial interpretation may lead a particular subject to the following conclusion, reading the conclusion off from c to a :




A test of the initial representation establishes its invalidity since both paths now become positive:



The really prudent subject, however, would do well not simply to rely on this test, but also to try to construct a conclusion in the converse direction.

The initial representation above suggests: No A are C, but the test modifies this to:



Try as one will, this conclusion cannot be falsified. Thus, the correct answer may only be obtained after a succession of tests.

The Logical Status of the Analogical Theory

The heuristic process embodied in the analogical theory is plainly outside logic. However, the test phase introduces a logical assessment of the validity of putative conclusions: If the tests are properly carried out, then any conclusion that remains will invariably be a logically valid one. Hence, the essence of the theory is that subjects are very good at drawing conclusions on heuristic grounds but generally less efficient at submitting them to logical tests.

As part of the process of developing the analogical theory, a number of different versions of it were modelled in the form of computer programs in a list-processing language, POP-2. The final model implemented a number of simplifications of the theory purely for ease of computation. However, quite unexpectedly, these modifications threw some light on the status of the theory and on the origins of the logical theory of syllogisms.

In specifying the tests to be carried out on initial representations it proved to be convenient to consider syllogisms in three basic categories: those with affirmative premises, those with an affirmative and a negative premise, and those with negative premises. With affirmative premises, it turns out that whenever one path can be broken, all of them can be broken. There was an extremely simple implementation of this principle: The relevant procedure simply looked for a middle item that was not linked to any end items, and whenever such an item was found the program indicated that no valid conclusion could be drawn. This procedure sacrifices psychological plausibility for the sake of simplicity: It cuts out a whole series of redundant processes that are likely to occur when logically naive individuals reason. However, this abstraction from actual behavior corresponds directly to one of the traditional laws of the syllogism: The middle term must be distributed at least once in a valid syllogism (see Cohen & Nagel, 1934, p. 79). With an affirmative and a negative premise, such as "Some A are B" and "No C are B," there is a similar shortcut. A prudent subject ought to test one conclusion and then, if it is invalid, test its converse. Rather than go through the whole complicated procedure illustrated above, the program merely checked

whether any term distributed in a conclusion was also distributed in the representation of the premises. If there was no such correspondence, the conclusion was invalid. This procedure also corresponds to a traditional rule for syllogisms: No term may be distributed in the conclusion which is not distributed in the premises. With negative premises, it was also unnecessary to specify a process of forming links between the end items. The principle implemented was simply to test for the presence of a path made up of two negative links since its existence is sufficient to establish the possibility of connecting the end items with a positive link. This shortcut amounts to the traditional rule of rejecting any conclusion drawn from two negative premises. Thus, we recovered all the major laws of the syllogism merely by simplifying the operation of the psychological principles. It may not be too farfetched to imagine that the original discoverers of those laws relied in part on analogous reflections on their own processes of thought.

AN EVALUATION OF THE ANALOGICAL THEORY

The Predictions of the Theory

The analogical theory was developed in order to account for the main sorts of response that were made in the two experiments. In fact, the theory predicts a total of 213 different responses to the 64 sorts of problems, an average of 3.3 responses per problem out of the nine possible responses, and the vast majority of responses that subjects made are within this set (92% of the first test and 95% of the second test in Experiment 2; see the tables in the Appendix). The theory predicts 23 responses that were not observed, but 16 of them were predicted to be relatively rare because they were incompatible with the figure of the premises.

In order to evaluate the theory, we shall consider initially the predictions that can be based on it about the relative difficulty of responding to premises correctly. The first basis from which such predictions derive is that if an initial representation is not tested logically a conclusion based on it may be erroneous. With some premises, the process of testing does not lead to any modifications: Such problems are predicted to be relatively easy. With other premises, the process of testing does lead to a modified representation, and hence a modified conclusion: Such problems are predicted to be relatively difficult. This difference was reliably confirmed by the results of the experiments. On the second test of Experiment 2, for example, 80.4% of responses to problems where a test leads to no modification were correct, whereas only 46.5% of responses to premises where a test leads to a modified conclusion were correct; the difference was apparent in the results of all 20 subjects.

The second basis from which predictions about difficulty can be derived is the figural bias created by the directional links. The bias predicts differences in accuracy within the set of problems that are unmodified by logical testing. The easiest of these problems will be those where the conclusions can be read off the representation in either direction (their converses are valid) and those problems where the conclusions can be read off in only one direction (their converses are invalid) but in accord with the figural bias. The hardest of these problems will be those where the conclusion can be read off in only one direction but with a figure that has no bias. The percentages of correct responses were in accord with this prediction: 88.1 and 85% versus 62.5%, respectively, on the second test of Experiment 2 (Sign test, $p < 0.035$). Figural bias also leads to a prediction about accuracy within the set of problems where the logical test leads to a modified conclusion. The easiest of these problems will be those where the conclusion can be read off in accord with the figural bias, slightly harder will be those problems where there is no figural bias, and the hardest will be those where the conclusion can be read off only in the direction opposite to the figural bias. The trend was in accord with this prediction: 73.3, 50.8, and 20.0% correct responses, respectively, on the second test of Experiment 2 (Page's $L = 266.5$, $p < 0.01$).

Comparable support for the theory is evident in the data for those premises that do not permit a valid conclusion to be drawn. It is plausible to assume that the easier it is to form paths the harder it will be to appreciate that there is no valid conclusion. Hence, premises without a figural bias will be easier than premises with a figural bias. The results of the experiments bear out the prediction: 78.2% correct responses to unbiased premises, and 64.8% correct responses to biased premises (a difference reflected in the performance of 18 out of the 20 subjects). Likewise, it should be easier to destroy an erroneous initial representation when there are fewer paths to be broken, that is, when premises are particular rather than universal. With affirmative premises, the percentages of correct responses in the second test were as follows: 82.5% where both premises were particular, 47.5% where one premise was particular, and 40% where neither premise was particular (Page's $L = 260$, $p < 0.05$). With one affirmative and one negative premise, the percentages of correct responses were 85% correct where both premises were particular, and 30% where one premise was particular (a difference reflected in the performance of 19 out of the 20 subjects). With negative premises, the percentages of correct responses were 93.8% where both premises were particular, 75.6% where one premise was particular, and 71.3% where neither premise was particular (Page's $L = 263$, $p < 0.01$). The relative ease of problems with two negative premises suggests that some subjects may have learned to interpret two negative links in a path as indeterminate, the shortcut implemented in the program.

The only wholly independent data available to test the predictions of the theory are results obtained by Mazzocco, Legrenzi, and Roncato (1974). These investigators required subjects to complete symbolic syllogisms by adding a second missing premise from a multiple choice of alternatives. For example, subjects were asked to complete the syllogism:

$$\begin{array}{c} \text{All A's are B's} \\ \dots \\ \hline \therefore \text{No A's are C's} \end{array}$$

As a matter of fact, 72% of the subjects selected "No B are C" and only 20% of the subjects selected "No C are B"; both answers are correct, but the theory predicts a figural bias toward the first one. This result is typical: As Mazzocco *et al.* report, where the given premise has the middle term as its predicate, 73% of the subjects selected a premise in which the middle term was the subject, thus creating a figure of the form $\frac{A-B}{B-C}$, but there was no such bias where the given premise has the middle term as its subject. Although providing a missing premise is rather different from drawing a conclusion from given premises, the analogical theory is readily extended to cope with it.

We assume that a subject represents the given premise, and then attempts to add to it links appropriate to form the path demanded by the conclusion. For example, with the problem

$$\begin{array}{c} \text{All A's are B's} \\ \dots \\ \hline \therefore \text{Some A's are not C's,} \end{array}$$

the first step is to represent the given premise

$$\begin{array}{ccc} \text{All A's are B's} & a & a \\ & \downarrow & \downarrow \\ & b & b \quad (b) \end{array}$$

and the second step is to construct the path(s) required by the conclusion:

$$\begin{array}{ccc} & a & a \\ & \downarrow & \downarrow \\ \therefore \text{Some A's are not C's} & b & b \quad (b) \\ & \perp & \downarrow \\ & c & (c) \end{array}$$

It is then necessary to formulate a premise corresponding to the new link(s): Some B's are not C's. However, the initial addition must be submitted to exactly the same process of logical testing that occurs in

ordinary syllogistic reasoning in order to ensure that the new premise guarantees the validity of the given conclusion. In the present case, the test consists in establishing that the negative path can be broken without doing violence to the premises:

All A's are B's	a	a	
	$\downarrow$	$\downarrow$	
Some B's are not C's	b	b	(b)
		$\downarrow$	$\perp$
		c	c

It is accordingly necessary to strengthen the negative pathways,

a	a	
$\downarrow$	$\downarrow$	
b	b	(b)
$\perp$	$\perp$	$\perp$
c	c	c ,

and this modification gives rise to the correct premise: *No B's are C's*. Problems of this sort in which an initial response is modified as a result of the logical test are predicted, of course, to be more difficult than those where the test has no effect on the initial response. In the present example, 32 out of the 50 subjects selected "Some B's are not C's" as the missing premise, and only two subjects made the correct selection of "*No B's are C's*." In general, there were 63.2% correct responses for the 10 problems unaffected by the logical test, 12.9% correct responses for the 17 problems where the test demands a modified response, and a minimal overlap between the two distributions.

There are, of course, other ways in which the analogical theory could be tested, and it is intended to investigate performance under time pressure in order to determine whether the conclusions that subjects draw correspond to those of the initial representations postulated by the theory. The final way in which we shall assess the theory is to compare it with other conjectures about syllogistic inference.

A Comparison of the Analogical Theory with Other Approaches

One of the most influential hypotheses about difficulties in syllogistic inference concerns the mood of a syllogism. The so-called "atmosphere" hypothesis proposed by Woodworth and Sells (1935) and Sells (1936) suggests that people are predisposed to accept a conclusion that is congruent in mood with the premises. The theory has been succinctly formulated by Begg and Denny (1969): Whenever at least one premise is negative, the most frequently accepted conclusion will be negative; whenever at least one premise is particular, the most frequently accepted conclusion will likewise be particular; otherwise the bias is towards

affirmative and universal conclusions. Revlis (1975a,b) has developed an information-processing model that allows errors to occur in working out the joint atmosphere of the two premises. It also assumes that if the atmosphere of the premises is incongruent with a given conclusion, a subject responds that the syllogism is invalid, or else considers the next possible conclusion in a multiple-choice test. The heuristic stage of the analogical theory yields initial conclusions that happen to be largely in accord with the atmosphere of the premises. However, the two theories diverge in at least four crucial ways.

First, the atmosphere hypothesis cannot even in principle explain the figural effect.

Second, the atmosphere hypothesis is unable to account for those conclusions that do not accord with the atmosphere of premises. For example, consider the results with the following premises:

All B are A	
No B are C	
∴ No A are C	(five subjects)
∴ No C are A	(three subjects)
∴ <i>Some A are not C</i>	(seven subjects)
∴ No valid conclusion	(four subjects)

The atmosphere hypothesis predicts only the results above the dotted line; if a principle of caution is introduced then it can be made to predict "*Some A are not C*" but only at the cost of also predicting "*Some C are not A.*" The four responses above are precisely those predicted by the analogical theory.

Third, when a subject draws a conclusion in his own words, then according to the atmosphere hypothesis he should never respond "No valid conclusion" because there is always a possible conclusion congruent with the atmosphere of the premises. Hence, the hypothesis is never able to explain the response, "No valid conclusion." The analogical theory predicts this response even in certain cases where a valid conclusion does exist (as in the example above): The results confirm that such responses are made.

Finally, it is well established that the apparent effects of atmosphere are invariably greater when a conclusion is valid than when it is invalid (see Sells, 1936; Revlis, 1975a,b; and the present results). The atmosphere hypothesis cannot, of course, explain such a phenomenon. In our view, it arises because most valid conclusions happen to be in accord with the atmosphere of the premises, and because subjects have recourse to an inferential mechanism that enables them both to make valid deductions and to refrain from drawing conclusions where none is warranted.

Another influential conjecture about the source of errors in syllogistic inference is Chapman and Chapman's (1959) theory of probabilistic inference. These authors argued that people often invalidly convert A and O statements; that is to say, "All A are B" is taken to imply "All B are A" and "Some A are not B" is taken to imply "Some B are not A." Such conversions can yield true conclusions in everyday life, and they could explain why the following syllogism is sometimes accepted as valid:

$$\begin{array}{l} \text{All A are B} \\ \text{All C are B} \\ \hline \therefore \text{All C are A} \end{array}$$

According to the Chapmans, subjects also assume on similar probabilistic grounds that entities with a predicate in common are likely to be the same sort of thing, for example,

$$\begin{array}{l} \text{Some A are B} \\ \text{Some C are B} \\ \hline \therefore \text{Some C are A,} \end{array}$$

and that entities that lack a common predicate are likely *not* to be the same sort of thing, for example,

$$\begin{array}{l} \text{Some A are B} \\ \text{Some C are not B} \\ \hline \therefore \text{Some C are not A.} \end{array}$$

Since the Chapmans investigated only invalid syllogisms, we cannot be entirely sure what their predictions would be for premises that permit a valid deduction. However, they do make clear predictions for 44 of our problems, comprising 37 premise pairs that permit no valid conclusions and seven premise pairs that permit valid conclusions only in unorthodox figures. The predicted response was the most frequent one in our data for only seven out of the 44 problems. Hence, their own corroboration of the theory may largely depend upon the use of symbolic materials and a multiple-choice test limited to conclusions in orthodox figures. A more explicit version of the illicit conversion hypothesis has recently been proposed by Revlis (1975a,b). Like his model based on the atmosphere effect, this model involves some assumptions additional to those proposed by the original theorists. However, even in this revised form, the conjecture does not receive very convincing support from the results of Revlis's or our experiment.

However, there is little doubt that subjects do sometimes argue from an A or an O premise to its converse, particularly with symbolic

materials (Wilkins, 1928; Sells, 1936). Why should this be so? We assume that such errors arise from the forgetting of the optional unlinked elements in a representation (or even perhaps a failure to include them in the first place). Thus, an A premise of the form "All A are B" may be erroneously represented as

$$\begin{array}{cc} a & a \\ \downarrow & \downarrow \\ b & b \end{array}$$

from which the converse, "All B are A," can be derived. Likewise, the erroneous representation of "Some A are not B" readily yields the converse assertions. However, the fact that subjects will accept the converse of an A or an O premise as valid provides no direct evidence for a process of conversion, licit or illicit, in syllogistic inference. Moreover, if such a process readily occurred, it would eliminate the figural effect. A figure of the form $\frac{A-B}{B-C}$ should be just as likely to yield a conclusion of the form $C-A$ as one of the form $A-C$. In developing the computer model of the analogical theory, we did at one time devise a program in which optional unlinked elements were "forgotten"; while such failures may indeed be one source of human error, they will not alone suffice to explain the experimental results. One decisive reason that led us to abandon this variant model is that it predicts far too many errors that do not occur. For example, with the premises

All A are B
Some C are not B,

it predicts the fallacious conclusion, No C are A. This and other similar errors for a variety of moods are seldom made.

In the case of those items that the Chapmans regarded as crucial tests between their account and the atmosphere predictions, our data fail to substantiate either approach decisively. Consider, for instance, the fate of the respective predictions for premises in the EI and IE moods, for example,

No B are A
Some C are B

Atmosphere	∴ Some A are not C	(no subjects)
predictions	∴ <i>Some C are not A</i>	(14 subjects)

Probabilistic	∴ No A are C	(no subjects)
predictions	∴ No C are A	(2 subjects).

Evidently, the atmosphere theory fails to explain the bias toward *Some C are not A*, and the probabilistic theory predicts only a small minority of

responses. The analogical theory, however, predicts two main responses: No *C* are *A* and *Some C are not A*. It appears to offer a more powerful explanation of syllogistic inference than either of these earlier theories.

The first explicit attempt to specify the mental processes involved in syllogistic inference is Erickson's (1974) set-theoretic model. According to Erickson, the premises are represented in forms equivalent to Euler circles. Thus, the representation of "All *A* are *B*" involves two separate mental diagrams: a circle representing set *A* included within a circle representing set *B*, and, since the two sets may be coextensive, a circle representing set *A* coincident with a circle representing set *B*. The representation of "Some *A* are *B*" requires four separate mental diagrams: set *A* overlapping set *B*, set *B* included in set *A*, set *A* included in set *B*, and set *A* co-extensive with set *B*. It is, of course, unlikely that subjects will be careful enough to consider all these possibilities, and Erickson assumes that such a failing is one source of error in inference. In particular, he assumes that "All *A* are *B*" is often treated as simply denoting that set *A* and set *B* are co-extensive. Likewise, since the combination of the representations of premises can present a considerable combinatorial problem, Erickson considers an alternative hypothesis in which subjects construct only one actual combination of representations selected at random from the total number of possibilities (each of which is assumed to be equiprobable). This procedure will invariably come up with a conclusion, and hence the model cannot predict a response of "No valid conclusion." When subjects formulate a conclusion to characterize the results of this process, they are supposed to select statements that agree with the mood of the premises; that is to say, Erickson assumes that the atmosphere effect operates at this stage. This assumption is necessary in order to account for the fact that, if not all combinations of premises are explicitly constructed, there will be occasions where a set overlap needs to be interpreted as "Some *A* are *C*" and other occasions where it needs to be interpreted as "Some *A* are not *C*."

A set-theoretic interpretation of quantified assertions has been explored by a number of psychologists (e.g., Johnson-Laird, 1970; Ceraso & Provitera, 1971; Neimark & Chapman, 1975). We find certain aspects of Erickson's (1974) model extremely plausible; indeed, the theory was put forward informally by Wason and Johnson-Laird (1972, p. 56-7). However, there are several phenomena that would seem to count against it and in favor of the analogical theory.

First, the fact that subjects readily respond "No valid conclusion" counts against Erickson's simple model in which only one combination of representations is constructed. The fact that they also make this response to premises that allow a valid conclusion counts against the full-scale model in which all combinations are constructed. Both these sorts of responses are, of course, predicted by the analogical theory.


Second, the figural effect obtained in the present experiments is a considerable embarrassment to set-theoretic representations such as Euler circles. They are symmetrical: The representation of "Some A are C" is identical to that of "Some C are A," and similarly the representation of "No A are C" is identical to that of "No C are A." The representations are quite without the directional component necessary to predict biases in the form of conclusions, and, at the very least, the theory would need to be supplemented in some way in order to account for the figural effect.

One such an assumption would be provided by Huttenlocher's finding of the importance of the grammatical subject in a variety of tasks (see, e.g., Huttenlocher and Weiner, 1971). Obviously, in the case of $\begin{smallmatrix} A & - & B \\ B & - & C \end{smallmatrix}$ premises only the A term can be maintained as the subject of the conclusion, in the case of $\begin{smallmatrix} B & - & A \\ C & - & B \end{smallmatrix}$ premises only the C term can be maintained as the subject of the conclusion, and in the case of the other two figures there can be no bias. Alternatively, the figural effect could be a consequence of the operations in working memory required to set up the initial representation of premises. In the spirit of Hunter's (1957) account of three-term series problems, we could argue that with $\begin{smallmatrix} A & - & B \\ B & - & C \end{smallmatrix}$ premises a subject encodes the first premise and can immediately add on to it a representation of the second premise: $A-B, \dots, B-C$, with a resulting bias toward a conclusion of the form $A-C$. It would be slightly harder to draw a conclusion from $\begin{smallmatrix} B & - & A \\ C & - & B \end{smallmatrix}$ premises because their middle terms are not adjacent, and so it would be necessary to recall the first premise to working memory to combine it with the second premise already there: $C-B, \dots, B-A$, with a resulting bias toward a conclusion of the form $C-A$. There are two alternative strategies for the remaining figures, and so there might well be no marked bias in their conclusions. These two explanations are not incompatible with each other. Although they could equally well be combined with the analogical or the set-theoretic theory, their addition to the latter is an *ad hoc* maneuver designed to save it from falsification, whereas their addition to the former is not really necessary.


Third, the set-theoretic theory lacks a suitable heuristic to make accurate predictions about performance when subjects draw their own conclusions from premises. Once again, we could suggest the following sort of principles: With affirmative premises, ensure that the intersections between sets are never empty, and with negative premises, ensure that the intersections between sets are always empty.

In short, it might appear that the analogical theory could be based on an Eulerian representation. However, the translation offers no gain in theoretical power and a considerable loss in flexibility. The analogical


representation can easily accommodate inferences involving particular individuals:




It can accommodate inferences involving quas numerical quantifiers:



What counts as a valid inference in this domain is more problematic because it has been relatively neglected by logicians (but see Altham, 1971). Yet it is clear that people readily make such inferences and that their knowledge of the world helps them to determine the relative sizes of classes. Thus, in the previous example the likely conclusion is, "Many fascists are dogmatic," whereas the following premises, superficially of the same form, are unlikely to elicit any conclusion: "Most geniuses are madmen," "Most madmen are in asylums." The relative sizes of the classes are entirely compatible with the representation:



The analogical representation also accommodates multiply quantified assertions:



This representation demands multiple links each denoting the appropriate relation, and it cannot be translated into Euler circles. Subjects who are logically naive run into some difficulty with such assertions, but they can interpret them and make inferences from them (Johnson-Laird, 1969a,b). The analogical theory accordingly provides a uniform method of representing quantified sentences, including ones that cannot be represented by Euler circles.

CONCLUSIONS

There are a variety of metaphors for human deductive reasoning. At one extreme, there is the idea in artificial intelligence of a uniform proof-procedure in which all deductions are handled by a single rule of inference applied to assertions in a standardized format. This approach may be intelligent, but it is extremely artificial. At the other extreme there is the idea of expressing every general assertion as a rule of inference couched in the form of a procedure. Such "theorems," most notably exploited in Carl Hewitt's problem-solving theory PLANNER (see Winograd, 1972) can even take into account information specific to their content, for example, hints as to how to achieve an inferential goal. There is also the intermediary notion of so-called "natural deduction" systems in logic in which a number of content-free rules of inference are introduced in a way that permits intuitiveness to take precedence over parsimony. There can be little doubt that human beings operate with both a logic resembling a natural deduction system and content-specific rules of inference of a sort postulated in PLANNER (see Johnson-Laird, 1975, for arguments in support of both suggestions). The theory of syllogistic inference that we have proposed differs in a number of respects from all of these approaches. It contains, of course, a simple heuristic for generating putative conclusions, a matter that falls outside the concerns of formal logic. The heuristic is epitomized by Forster's motto quoted at the head of this paper, though the connections that concerned him were of a different sort. Once a putative conclusion has been generated, it could be evaluated by a system with a single resolution rule of inference or a system with a variety of rules of inference. However, when we consider the testing procedure of the present model of syllogistic inference, it is not easy to classify it in these terms. It contains neither a single rule of inference nor a whole set of them: Rather, the rules are inherent in the way it models the entities and relations involved in a syllogism. Although it is hard to decompose this system into separate rules, its performance *can* be captured by a set of rules, the traditional rules of syllogistic inference. However, it would be misleading to think of it in these terms. We regard the difficulty, or at least lack of perspicuity, in matching it to rules of inference as an argument in favor of its psychological plausibility. If human reasoning followed a simple set of principles, the task of specifying them should have been solved long ago.

APPENDIX

The detailed predictions of the analogical theory are set out in Tables 6-9, together with the results of Experiment 2. The upper pair of statements in each cell are the theory's predictions about the initial conclusions. The statements below the dotted line correspond to the conclu-

sion, if any, forthcoming after the logical testing process. Thus, in the case of the problem “Some A are B” and “No C are B,” one finds the following results

No A are C	(6)	1
No C are A	(3)	7
<i>Some A are not C</i>	(3)	7
No valid conclusion	(5)	4 ,

where “*Some A are not C*” is the result of testing “No A are C,” and “No valid conclusion” is the result of testing “No C are A”. Numbers in parentheses correspond to the numbers of subjects making the response on the first test; other numbers are the results of the second test. Conclusions that are predicted to be relatively rare are included in parentheses; valid conclusions are italicized. In certain cases a representation is susceptible to more than one modification, for example,

		+	+
All A are B	<i>a</i>	<i>a</i>	
	↓	↓	
All C are B	<i>b</i>	<i>b</i>	(b)
	↑	↑	
	<i>c</i>	<i>c</i>	
	+	+	
			∴ All A are C (2) 2 subjects
			∴ All C are A (3) 5 subjects
		?	+
	<i>a</i>	<i>a</i>	
	↓	↓	
	<i>b</i>	<i>b</i>	(b)
	↑	↑	
	<i>c</i>	<i>c</i>	
	+	?	
			∴ Some A are C (4) 3 subjects
			∴ Some C are A (1) 1 subject
		?	?
	<i>a</i>	<i>a</i>	
	↓	↓	
	<i>b</i>	<i>b</i>	(b)
	↑	↑	
	<i>c</i>	<i>c</i>	
	?	?	
			∴ No valid conclusion: (9) 8 subjects

In such cases, the data within a cell are divided into three. The relatively few unpredicted responses that occurred with any frequency are preceded by an asterisk in the tables.

TABLE 6
THE FREQUENCIES OF THE CONCLUSIONS FOR A—B PREMISES IN EXPERIMENT 2, TOGETHER WITH
B—C
THE MEAN LATENCIES FOR THE CORRECT RESPONSES^a

		First premise				
Second premise	A	I	E	O		
A	All A are C (Some C are A) *Some A are C	(16) 17 (0) 1 (3) 1	(18) 16 (1) 3	No A are C (No C are A)	(9) 11 (5) 2	
				No valid conclusion (Some C are not A)	(5) 2 (1) 3	
	(3.8) 2.9	(4.7) 5.7			(9.5) 10.2	
I	Some A are C (Some C are A)	(13) 12 (0) 0	(7) 3 (0) 0	No A are C (No C are A)	(6) 3 (0) 2	
	No valid conclusion	(7) 8	(13) 16	No valid conclusion (Some C are not A) *Some A are C	(5) 4 (6) 8 (3) 0 (17.4) 10.3	
	(11.5) 4.4	(3.1) 1.9			(8.4) 7.1	
E	No A are C (No C are A)	(14) 17 (1) 1	(2) 2 (0) 0	No A are C (No C are A)	(10) 5 (0) 2	
				No valid conclusion	(9) 13	
	(9.0) 6.7	(7.1) 7.1			(16.4) 6.5	
O	Some A are not C (Some C are not A)	(12) 14 (0) 1	(10) 4 (0) 0	No A are C (No C are A)	(1) 0 (0) 0	
	No valid conclusion	(6) 5	(10) 16	Some A are not C (Some C are not A)	(7) 4 (1) 0	
	(9.2) 8.2	(5.2) 3.8		No valid conclusion	(9) 16 (14.5) 10.9	
					(7.7) 8.2	

^a Correct responses are italicized, and numbers in parentheses are from the first test and the other numbers are from the second test. Mean latencies are given in seconds. The results are set out in terms of the predictions of a model. The upper pair of statements in each cell corresponds to initial responses, and the statements below the dotted line correspond to the respective results, if any, of submitting the initial responses to the tests specified by the model. Responses that are predicted to be rare are included in parentheses. An asterisk indicates a response not predicted by the model. Only those such responses that occurred more than twice on a given test have been included, and hence the frequencies in a cell do not invariably sum to 20 (the number of subjects tested).

TABLE 7
THE FREQUENCIES OF THE CONCLUSIONS FOR B—A PREMISES IN EXPERIMENT 2, TOGETHER WITH
C—B
THE MEAN LATENCIES FOR THE CORRECT RESPONSES^a

First premise					
Second premise	A	I	E	O	
A	All C are A (Some A are C)	Some C are A (Some A are C)	No C are A (No A are C)	Some C are not A (Some A are not C)	(15) 15 (0) 3 (0) 0
	*Some C are A	No valid conclusion	*No valid conclusion	No valid conclusion	(1) 3 (13.7) 10.1
	(5.7) 5.7	(4.4) 6.6	(7.2) 7.2		(3) 3
I	Some C are A (Some A are C)	Some C are A (Some A are C)	No C are A (No A are C)	Some C are not A (Some A are not C)	(15) 18 (3) 0 (0) 1
	*No valid conclusion	No valid conclusion	Some C are not A (No valid conclusion)	No valid conclusion	(9) 15 (10) 14 (5) 3
	(6.3) 6.1	(5.5) 6.0	(9.8) 8.9	(7.6) 5.3	(8) 3 (0) 1 (9) 16
E	No C are A (No A are C)	No C are A (No A are C)	No C are A (No A are C)	No C are A (No A are C)	(1) 1 (0) 0 (0) 0
	No valid conclusion (Some A are not C)	No valid conclusion (Some A are not C)	No valid conclusion	Some C are not A (Some A are not C)	(4) 2 (0) 0 (0) 0
	*Some C are not A	*Some C are not A		No valid conclusion	(13) 17 (15.5) 13.2
O	(—) (—)	(38.8) 11.7	(11.9) 8.6	Some C are not A (Some A are not A)	(0) 3 (2) 1 (17) 16
	Some C are not A (Some A are not C)	Some C are not A (Some A are not C)	No C are A (No A are C)	No valid conclusion	(0) 0 (0) 0 (0) 0
	No valid conclusion	No valid conclusion	Some C are not A (Some A are not C)	No valid conclusion	(16) 18 (12.6) 7.3 (7.7) 4.6

^a See footnote a of Table 6 for explanation of presentation of data.

TABLE 8
THE FREQUENCIES OF THE CONCLUSIONS FOR A—B PREMISES IN EXPERIMENT 2, TOGETHER WITH
MEAN LATENCIES FOR THE CORRECT RESPONSES^a

Second premise	First premise				
	A	I	E	O	
A	All A are C (2) 2	Some A are C (4) 5	No A are C (9) 9	Some A are not C (7) 5	
	All C are A (3) 5	Some C are A (2) 5	No C are A (11) 6	Some C are not A (2) 6	
	Some A are C (4) 3	No valid conclusion (13) 9	*Some C are not A (0) 3	No valid conclusion (7) 9	
	Some C are A (1) 1				
I	No valid conclusion (7.9) 3.6				
	Some A are C (4) 1	(7.2) 13.5	(9.6) 7.3	(11.7) 8.9	
	Some C are A (5) 4	Some A are C (4) 3	No A are C (2) 1	Some A are not C (2) 2	
	No valid conclusion (11) 15	Some C are A (1) 0	No C are A (4) 1	Some C are not A (1) 3	
E		No valid conclusion (15) 17	No valid conclusion (6) 4	No valid conclusion (16) 14	
	(12.1) 6.6		Some C are not A (7) 13		
	No A are C (7) 3	(3.6) 4.2	(8.0) 8.4	(9.5) 7.0	
	No C are A (10) 12	No A are C (6) 1	No A are C (2) 2	No A are C (1) 0	
O			No C are A (0) 0	No C are A (1) 2	
	(7.7) 9.7		No valid conclusion (18) 17	Some A are not C (2) 0	
	Some A are not C (3) 0	(17.1) 10.8	(9.5) 4.5	Some C are not A (1) 3	
	Some C are not A (9) 14	Some A are not C (2) 0	No A are C (1) 0	No valid conclusion (14) 14	
	No valid conclusion (7) 6	No valid conclusion (14) 18	Some A are not C (2) 0	Some A are not C (1) 0	
			Some C are not A (1) 2	Some C are not A (1) 0	
			No valid conclusion (13) 17	No valid conclusion (18) 20	
	(9.2) 8.2	(9.6) 6.4	(10.8) 11.5	(6.6) 5.8	

^a See footnote a of Table 6 for explanation of presentation of data.

TABLE 9
THE FREQUENCIES OF THE CONCLUSIONS FOR B—A PREMISES IN EXPERIMENT 2, TOGETHER WITH
B—C
THE MEAN LATENCIES FOR THE CORRECT RESPONSES^a

Second premise		First premise			
	A	I	E	O	
A	All A are C (8) 1	Some A are C (11) 8	No A are C (4) 3	Some A are not C (0) 1	
	All C are A (2) 3	Some C are A (7) 12	No C are A (4) 4	Some C are not A (12) 14	
	Some A are C (4) 8		No valid conclusion (6) 4	No valid conclusion (3) 5	
	Some C are A (1) 5		Some C are not A (5) 7	*Some A are C (3) 0	
I	*No valid conclusion (9.0) 8.5	(6.4) 5.1	(9.6) 8.3	(9.5) 9.5	
	Some A are C (12) 11	Some A are C (3) 2	No A are C (1) 2	Some A are not C (2) 0	
	Some C are A (7) 9	Some C are A (0) 0	No C are A (3) 0	Some C are not A (5) 1	
		No valid conclusion (17) 18	No valid conclusion (8) 4	No valid conclusion (13) 19	
E	(8.5) 5.5	(3.6) 3.1	(14.2) 7.1	(5.9) 6.3	
	No A are C (9) 5	No A are C (2) 1	No A are C (2) 4	No A are C (0) 0	
	No C are A (0) 3	No C are A (0) 3	No C are A (1) 0	No C are A (1) 1	
	Some A are not C (4) 7	Some A are not C (9) 14	No valid conclusion (15) 13	Some A are not C (3) 1	
O	No valid conclusion (6) 4	No valid conclusion (8) 2	*Some A are not C (1) 3	Some C are not A (2) 3	
	(7.9) 11.5	(11.3) 8.1	(10.5) 7.0	No valid conclusion (11) 14	
	Some A are not C (13) 17	Some A are not C (5) 1	No A are C (0) 1	Some A are not C (0) 0	
	Some C are not A (1) 1	Some C are not A (0) 0	No C are A (0) 1	Some C are not A (1) 0	
	No valid conclusion (3) 1	No valid conclusion (14) 19	Some A are not C (3) 4	No valid conclusion (19) 20	
			Some C are not A (1) 0		
	(5.4) 7.9	(5.3) 5.5	No valid conclusion (14) 14	(5.2) 4.3	
			(9.6) 15.6		

^a See footnote a to Table 6 for explanation of presentation of data.

REFERENCES

- Altham, J. E. J. *The logic of plurality*. London: Methuen, 1971.
- Begg, I., & Denny, J. P. Empirical reconciliation of atmosphere and conversion interpretations of syllogistic reasoning errors. *Journal of Experimental Psychology*, 1969, **81**, 351–354.
- Ceraso, J., & Provitera, A. Sources of error in syllogistic reasoning. *Cognitive Psychology*, 1971, **2**, 400–410.
- Chapman, I. J., & Chapman, J. P. Atmosphere effect re-examined. *Journal of Experimental Psychology*, 1959, **58**, 220–226.
- Cohen, M. R., & Nagel, E. *An introduction to logic and scientific method*. London: Routledge & Kegan Paul, 1934.
- Erickson, J. R. A set analysis theory of behavior in formal syllogistic reasoning tasks. In R. Solso (Ed.), *Theories in cognitive psychology: The Loyola Symposium*. Potomac, MD: Erlbaum, 1974.
- Frase, L. T. Associative factors in syllogistic reasoning. *Journal of Experimental Psychology*, 1968, **76**, 407–412.
- Hunter, I. M. L. The solving of three-term series problems. *British Journal of Psychology*, 1957, **48**, 286–298.
- Huttenlocher, J., & Weiner, S. Comprehension of instructions in varying contexts. *Cognitive Psychology*, 1971, **2**, 369–385.
- Johnson-Laird, P. N. On understanding logically complex sentences. *Quarterly Journal of Experimental Psychology*, 1969a, **21**, 1–13.
- Johnson-Laird, P. N. Reasoning with ambiguous sentences. *British Journal of Psychology*, 1969b, **60**, 17–23.
- Johnson-Laird, P. N. The interpretation of quantified sentences. In G. B. Flores D'Arcais & W. J. M. Levelt (Eds.), *Advances in psycholinguistics*. Amsterdam: North-Holland, 1970.
- Johnson-Laird, P. N. Models of deduction. In R. J. Falmagne (Ed.), *Reasoning: Representation and process in children and adults*. Hillsdale, NJ: Erlbaum, 1975.
- Mazzocco, A., Legrenzi, P., & Roncato, S. Syllogistic inference: The failure of the atmosphere effect and the conversion hypothesis. *Italian Journal of Psychology*, 1974, **2**, 157–172.
- Neimark, E. D., & Chapman, R. H. Development of the comprehension of logical quantifiers. In R. J. Falmagne (Ed.), *Reasoning: Representation and process in children and adults*. Hillsdale, NJ: Erlbaum, 1975.
- Revlis, R. Two models of syllogistic reasoning: Feature selection and conversion. *Journal of Verbal Learning and Verbal Behavior*, 1975a, **14**, 180–195.
- Revlis, R. Syllogistic reasoning: Logical decisions from a complex data base. In R. J. Falmagne (Ed.), *Reasoning: Representation and process in children and adults*. Hillsdale, NJ: Erlbaum, 1975b.
- Sells, S. B. The atmosphere effect: An experimental study of reasoning. *Archives of Psychology*, 1936, **29**, 3–72.
- Wason, P. C., & Johnson-Laird, P. N. *Psychology of reasoning: Structure and content*. Cambridge, MA: Harvard University Press; London: Batsford, 1972.
- Wilkins, M. C. The effect of changed material on the ability to do formal syllogistic reasoning. *Archives of Psychology*, 1928, **16**, No. 102.
- Winograd, T. *Understanding natural language*. New York: Academic Press, 1972.
- Woodworth, R. S. *Experimental psychology*. New York: Holt, 1938.
- Woodworth, R. S., & Sells, S. B. An atmosphere effect in formal syllogistic reasoning. *Journal of Experimental Psychology*, 1935, **18**, 451–460.