

On the conflict between logic and belief in syllogistic reasoning

J. St. B. T. EVANS, JULIE L. BARSTON, and PAUL POLLARD
Plymouth Polytechnic, Plymouth PL4 8AA, England

Three experiments are reported that investigate the weighting attached to logic and belief in syllogistic reasoning. Substantial belief biases were observed despite controls for possible conversions of the premises. Equally substantial effects of logic were observed despite controls for two possible response biases. A consistent interaction between belief and logic was also recorded; belief bias was more marked on invalid than on valid syllogisms. In all experiments, verbal protocols were recorded and analyzed. These protocols are interpreted in some cases as providing rationalizations for prejudiced decisions and, in other cases, as reflecting a genuine process of premise to conclusion reasoning. In the latter cases, belief bias was minimal but still present. Similarly, even subjects who focus primarily on the conclusion are influenced to an extent by the logic. Thus a conflict between logic and belief is observed throughout, but at several levels of extent.

An important debate in cognitive psychology surrounds the notion of rationality with respect to human inference (see Cohen, 1981, and associated commentaries). Recent reviews by Evans (1982) and Nisbett and Ross (1980) have stressed the role of apparently irrational processes in the study of inductive and deductive inference, respectively (see also Pollard, 1982). However, theories favoring a rationalist interpretation of inferential behavior still hold a dominant position in the recent literature (see, for example, the collections of papers edited by Falmagne, 1975, and Revlin & Mayer, 1978). In the case of deductive reasoning, much of the argument centers on the use by subjects of a system of logic, whether of the philosopher's variety (cf. Henle, 1962) or of an alternative "natural" type (e.g., Braine, 1978).

The nonlogical or antirational position is sometimes misinterpreted as denigrating man's proven intelligence. What is in fact suggested is an alternative conception of that intelligence. The "rationalist" position entails the supposition that the reasoner proceeds by forming an abstract representation of problem information and applying a general set of inferential rules to its logical structure, regardless of its content. This notion is clearly embodied, for example, in Piaget's theory of formal operations (cf. Inhelder & Piaget, 1958). The alternative argument stressed here is that specific features of problem content, and their semantic associations, constitute the dominant influence on thought (see Evans, 1982, for extended discussion).

The authors would like to acknowledge the Social Science Research Council for financial support of the research reported in this paper. The manuscript was prepared for publication while the first author was on sabbatical leave at the University of Florida. Requests for reprints should be sent to J. St. B. T. Evans, Department of Psychology, Plymouth Polytechnic, Plymouth PL4 8AA, England.

In this paper, we will focus on the alleged "belief-bias" effect in reasoning. The claim is that when presented with deductive arguments to evaluate, subjects will make judgments upon a priori beliefs rather than on the basis of logical argument. Specifically, they will tend to endorse arguments whose conclusions they believe and reject arguments whose conclusions they disbelieve, irrespective of their actual validity. A number of authors have claimed evidence of such an effect in syllogistic reasoning (e.g., Feather, 1964; Gorden, 1953; Henle & Michael, 1956; Janis & Frick, 1943; Janis & Terwilliger, 1962; Kaufman & Goldstein, 1967; Lefford, 1946; Morgan & Morton, 1944; Wilkins, 1928; Wilson, 1965; Winthrop, 1946).

Most of these studies are, however, open to criticism on a variety of grounds (cf. Evans, 1982). Some find only weak effects, not backed by tests of statistical significance; others use poorly worded problems, fail to control for other factors that influence reasoning, or risk carryover effects by having subjects rate the believability and validity of arguments in the same session. Revlin and Leirer (1978) and Revlin, Leirer, Yopp, and Yopp (1980) have raised other problems, such as a failure to control for the pragmatic convertibility of the logical premises, that may affect the logical status of the problem representation. Revlin et al. argue that a rational reasoner, as proposed by the model of Revlin (1975a, 1975b), should not manifest belief biases, and they rightly contend that the aging literature on the subject should be opened up to investigation with improved methodologies. While finding some evidence of belief-bias effects, Revlin et al. argue that these are relatively weak in comparison with the logical tendencies observed, when due allowance is made for personalized representations of the premises. The present study questions the accuracy of this conclusion.

First, we must briefly review the claims of the Revlin

(1975a, 1975b) model, more specifically known as the conversion model of syllogistic reasoning. A classical syllogism consists of two premises and a conclusion relating three categories. Valid syllogisms are those whose premises logically determine their conclusions. The structure of syllogisms is described in detail below.

Revlin's (1975a, 1975b) model directs our attention to the personal encodings of the premises of the syllogism. It is proposed that when the premises contain information about categories that are already available to the reasoner, long-term memory may provide working memory with more information than may have been contained or intended in the presented material. Hence, the reasoner makes his or her judgments based upon "too much" information, and not only on the specific content of the problem. Revlin suggests that this is manifested in terms of the kinds of immediate inferences that the reasoners are willing to make when presented with a proposition. For example, it is claimed that when shown the abstract relations "All A are B," reasoners often infer that "All B are A." However, due to pragmatic implications, conversion may be blocked in some concrete relations (e.g., "All dogs are animals"). It is this kind of use of personal knowledge that the conversion model claims will affect the validity judgments on categorical syllogisms. It is proposed that subjects will be correct in their judgments in either of two conditions: (1) when the premises of the problem are converted, but the logical conclusion is fortuitously the same in the converted and the original forms of the problem, and (2) when the subject's knowledge of the world blocks illicit conversion.

Revlin et al. (1980; see also Revlin & Leirer, 1978) report evidence to support these hypotheses. As mentioned earlier, they also find significant, although they claim "weak," evidence of belief bias when conversion of the premises is fully controlled. However, there are several methodological problems identified by Pollard (1979) that may have led Revlin et al. to underestimate the extent of the belief-bias effect. For example, Revlin and Leirer (1978) claim that belief and logic conflict in the following problem: "No U.S. governors are members of the Harem Club. Some Arabian sheiks are members of the Harem Club. Therefore: (a) All Arabian sheiks are U.S. governors. (b) No Arabian sheiks are U.S. governors. (c) Some Arabian sheiks are U.S. governors. (d) Some Arabian sheiks are not U.S. governors. (e) None of the above is proved."

The "believable" answer is claimed to be Answer b, but Answer d is also empirically true. Since Answer d is the logically correct answer as well, Revlin and Leirer's (1978) claim that subjects are overriding personal beliefs in choosing it is doubtful. This problem arises as a result of the multiple-choice technique, and it can be avoided by presenting only one conclusion for evaluation, as in the experiments to be reported here.

A second problem is that the Revlin (1975a, 1975b) studies employed primarily valid syllogisms. There is

evidence to suggest that belief-bias effects may be more marked on indeterminate syllogisms, that is, on those whose conclusions do not follow logically (Kaufman & Goldstein, 1967).

A third problem is that Revlin and Leirer (1978) did not control for the effects of "atmosphere," an alleged bias to choose syllogistic conclusions that share syntactic features with the premises, regardless of logical validity. Although existence of this effect, proposed by Woodworth and Sells (1935), is regarded primarily as an artifact of conversion by Revlin (1975a, 1975b), there is much evidence compatible with the hypothesis (see Evans, 1982). Since Gorden (1953) has specifically suggested that belief bias is weaker than atmosphere, it is advisable to control for its possible effects. In fact, the logically correct answer to Revlin et al.'s (1980) valid syllogisms was also the conclusion favored by atmosphere, which may have led Revlin et al. to overestimate subjects' logical abilities.

Experiment 1 was designed to test the relative weighting given to logic and belief in syllogistic reasoning, controlling for conversion of premises as Revlin (1975a, 1975b) requires, but also improving upon his methodology in the three respects outlined above. In addition, subjects were asked to provide retrospective verbal justifications of their decisions. On a different reasoning task, Wason and Evans (1975) claimed that such reports indicated no evidence of insight into the origin of responses apparently induced by a nonlogical bias and produced logical sounding post hoc rationalizations. Further investigation of this phenomenon is of interest with respect to the recent debate about the interpretation of verbal protocols (see, for example, Ericsson & Simon, 1980; Nisbett & Wilson, 1977).

EXPERIMENT 1

The type of problems used in Experiment 1 was categorical syllogisms. A syllogism is a deductive argument consisting of two premises and a conclusion. The two premises make statements about the relations between three terms: a major term (P), a minor term (S), and a middle term (M).

The figure of a syllogism indicates the position of the middle term in the premises. There are four possible figures, shown in Figures 1-4.

Syllogisms are composed of a combination of four basic types of statement: (1) The universal affirmative statement: All X are Y (symbolized by "A"). (2) The particular affirmative statement: Some X are Y (symbolized by "I"). (3) The universal negative statement: No X are Y (symbolized by "E"). (4) The particular negative statement: Some X are not Y (symbolized by "O"). The types of statement that occur in any particular syllogism specify its mood.

The form of a syllogism may be completely described, therefore, by stating its mood and figure. Thus a (valid) EIO-2 syllogism is of the following form: "No A are B.

M	-	P
S	-	M
<hr/>		
S	-	P

Figure 1.

P	-	M
S	-	M
<hr/>		
S	-	P

Figure 2.

M	-	P
M	-	S
<hr/>		
S	-	P

Figure 3.

P	-	M
M	-	S
<hr/>		
S	-	P

Figure 4.

Some C are B. Therefore, some C are not A." This was one of the types of syllogism used in Experiment 1.

Both E and I statements are "legally" convertible, which means that the terms of the statement may be reversed without altering its meaning in logic. Since the two premises of the syllogisms used in this experiment are of this type, even if subjects do in fact convert premises, as Revlin and his associates suggest, this could not in itself be the cause of logical errors. The conversion model, therefore, predicts that reasoning on this task will be logical, regardless of the type of material used, and there should be no belief bias.

In the Introduction, a reference was made to atmosphere bias. This is a nonlogical bias first investigated by Woodworth and Sells (1935; see also Begg & Denny, 1969). Briefly, this effect is due to the type of quantifiers used in the two premises, which combine to create an "atmosphere," which predisposes subjects to accept a conclusion containing specific quantifiers. Since the

conclusion to EI premises favored by atmosphere bias is an O conclusion, all conclusions are favored by the bias and it cannot confound any comparisons between problems.

Method

Materials. Half of the syllogisms presented were valid, and half were invalid (i.e., their conclusions did not follow logically from the premises). The following two syllogisms were used throughout: "No A are B. Some C are B. Therefore, some C are not A" (valid). "No A are B. Some C are B. Therefore, some A are not C" (invalid).

Note that invalid conclusions were thus of the form C-A and invalid, A-C. Both syllogisms are in Figure 2, but the latter reverses the traditional premise order, so that the mood remains the same for both valid and invalid syllogisms.

The materials were chosen so that the conclusions of the syllogisms would appear "true" when the terms were presented in one order, but "false" when the order of terms was reversed. The experimenters' intuitions were checked by having a group of 32 subjects, who did not participate in the experiment, rate the conclusions for believability (see Table 1). It will be seen that the differences in ratings between "true" and "false" sentences are very marked.

Half of the valid conclusions presented to subjects were believable, and half were unbelievable. For example, the following syllogism, which is valid, has a believable conclusion: "No cigarettes are inexpensive. Some addictive things are not cigarettes." This valid syllogism, on the other hand, has an unbelievable conclusion: "No addictive things are inexpensive. Some cigarettes are inexpensive. Therefore, some cigarettes are not addictive."

For invalid syllogisms, as for valid, half were believable, and half were not: "No addictive things are inexpensive. Some cigarettes are inexpensive. Therefore, some addictive things are

Table 1
Believability Ratings of Conclusions Used in the Three Experiments

	Materials	Mean	SD
A (Experiments 1, 2, and 3)			
T	Some highly trained dogs are not police dogs	6.44	.89
F	Some police dogs are not highly trained	2.75	1.84
T	Some nutritional things are not vitamin tablets	5.75	2.11
F	Some vitamin tablets are not nutritional things	3.81	1.64
T	Some addictive things are not cigarettes	6.25	1.88
F	Some cigarettes are not addictive things	2.81	1.64
T	Some rich people are not millionaires	5.94	1.57
F	Some millionaires are not rich people	3.00	1.90
B (Experiment 3)			
T	Some religious people are not priests	6.94	.25
F	Some priests are not religious people	1.69	1.14
T	Some healthy people are not astronauts	6.94	.25
F	Some astronauts are not healthy people	3.75	2.18
T	Some good swimmers are not deep sea divers	6.31	1.08
F	Some deep sea divers are not good swimmers	2.75	2.18
T	Some well educated people are not judges	6.44	1.55
F	Some judges are not well educated people	3.31	1.85

Note—All items were rated on a 7-point scale from 1 = completely unbelievable to 7 = completely believable. Materials A were rated by two groups of 16 subjects and Materials B by two separate groups of 16 subjects. Each subject rated four statements, one from each context, of which two were "true" and two "false."

not cigarettes" (invalid, believable conclusion). "No cigarettes are inexpensive. Some addictive things are inexpensive. Therefore, some cigarettes are not addictive" (invalid, unbelievable conclusion).

There were thus four types of problem: valid conclusion, believable or unbelievable, and invalid conclusion, believable or unbelievable. Of course, the four problems actually given to each subject used different problem contents.

To reduce the artificiality of the task, problems were presented in the form of prose passages that were approximately 80 words in length. Four different types of passage content were used, each taking the form of a current affairs article. The four topics were: (1) public response to the behavior of police dogs, (2) the provision of aid for third-world countries, (3) attempts to reduce the number of people smoking cigarettes, and (4) the relationship between wealth and hard work. The following is an example of Passage Type 1: "Dogs are used extensively for the purpose of guarding property, guiding the blind and so on. No highly trained dogs are vicious. However, many people believe that their temperament cannot be trusted. The police service use dogs a great deal in their work. Some police dogs are vicious and although fatal accidents are rare, there is still growing concern over their widespread use." "If the above passage is true, does it follow that: Some highly trained dogs are not police dogs?" (This conclusion is invalid, but believable.)

Design. Each subject received each of the four types of prose passages and problem types, solving four problems in all. Combination of problem type and passage type was balanced in a Latin square design, and presentation order was randomized.

Subjects. Twenty-four undergraduates at Plymouth Polytechnic acted as paid volunteers. They had no previous experience of syllogistic reasoning tasks and were tested individually.

Procedure. Task and instructions. The instructions and problems for each subject were presented on typed cards. All problems were presented individually, and each problem card remained in front of subjects for reference when decisions were explained. The instructions began as follows:

"This is an experiment to test people's reasoning ability. You will be given four problems. In each case, you will be given a prose passage to read and asked if a certain conclusion may be logically deduced from it. You should answer this question on the assumption that all the information given in the passage is, in fact, true. If you judge that the conclusion necessarily follows from the statements in the passage, you should answer 'yes,' otherwise 'no.'

"Please take your time and be sure that you have the right answer before stating it. When you have decided, I will then ask you to explain why you believe the conclusion to be valid or invalid as the case may be. Any questions?"

Subjects' protocols were recorded on a tape recorder for later analysis.

Protocol scoring. Each protocol was scored on a yes/no basis on two criteria: (1) presence or absence of a reference to both the logically relevant premises, and (2) presence or absence of references to irrelevant information, either within the passage or extraneous.

Results and Discussion

The percentage frequencies of subjects accepting the conclusion (i.e., deeming the argument to be valid) are shown for each type of problem in Table 2. As predicted, there was a substantial effect of "belief bias" (i.e., a tendency over all problems to accept more believable than unbelievable conclusions; $p < .01$, one-tailed sign test). There was also a tendency to accept more valid than invalid arguments overall ($p < .02$, one-tailed) and a significant Belief by Validity interaction ($p < .05$). The nature of this interaction accords

Table 2
Percentage Frequency of Subjects Accepting
Conclusions in Experiment 1 (n = 24)

	Believable	Unbelievable
Valid	92	46
Invalid	92	8

with the findings of Kaufman and Goldstein (1967) that belief bias is more marked for invalid than for valid syllogisms.

These results cannot be reconciled with the essentially rationalist approach of Revlin and Leirer (1978) and Revlin et al. (1980). Their claim that previous evidence of belief bias could be an artifact of uncontrolled premise conversion cannot be applied to the present experiment, in which only E and I premises were involved. Their further suggestion, on the basis of their own data, that belief biases are weak in comparison with rational processes is also inconsistent with our results. The suggestion in the introduction that their methodology led to an underestimate of the true extent of belief bias is strongly confirmed in our data. They found subjects to be correct when logic accorded with belief on 83% of occasions, and when logic conflicted with belief, on more than 67% of occasions. The corresponding percentages in Experiment 1 were 92% and 27%.

The present results do accord well with the Evans (1982) two-factor theory, which claims that reasoning responses reflect a competition between logical and nonlogical tendencies. This theory has previously been applied mostly to conditional reasoning problems (e.g., Evans, 1977a, 1977b). Related to this is the dual process theory of Wason and Evans (1975), who claim that the verbalizations observed on their reasoning problems reflect primarily a type of thought different from that determining the reasoning response. Specifically, they found that subjects tended to rationalize responses attributed to nonlogical biases. Similar trends should be found in the protocols collected in Experiment 1.

However, Ericsson and Simon (1980) have argued persuasively that protocols may reveal the locus of the subjects' attention, or the information heeded by the problem solver. It could be that the Wason and Evans' (1975) rationalizations were due to their asking the subjects to justify the responses given. The important issue in the present study is that of whether subjects base their reasoning on the logical premises or on extraneous beliefs. The protocols were consequently scored separately for the presence or absence of references to both, and the results are shown in Table 3.

Of particular interest are the two conditions in which logic and belief conflict. If subjects are rationalizing, then we might expect that their protocol ratings would interact with their response to the problem. That is, subjects who give the logical response should make more references to the premises and those favoring belief

Table 3
Classification Frequencies for Protocols of Experiment 1 on the
Two Criteria, Broken Down by Response Given (n = 24)

	Both Premises		Irrelevant Information	
	Response Yes	Response No	Response Yes	Response No
Valid-Believable				
M	<i>9</i>	1	9	1
NM	<i>13</i>	1	<i>13</i>	1
Valid-Unbelievable				
M	6	2	2	8
NM	5	11	9	5
Invalid-Believable				
M	6	<i>1</i>	14	<i>1</i>
NM	16	<i>1</i>	8	<i>1</i>
Invalid-Unbelievable				
M	1	2	1	<i>13</i>
NM	1	20	1	9

Note—Correct responses italicized. M = mentioned, NM = not mentioned.

should make more references to irrelevant information. The same prediction would also be made if the protocols were assumed to reflect the actual information on which the subjects' reasoning was based (an attempt will be made to distinguish these possibilities in the general discussion). Only the valid-unbelievable condition produced a sufficiently even split of "yes" and "no" responses to permit test of this hypothesis. The predicted interaction was present and significant for references to irrelevant information ($p = .026$, one-tailed Fisher exact probability test), but not for references to the logical premises. It is also relevant to note that in the invalid-believable condition, subjects who accepted the conclusion (in accordance with beliefs) showed the highest ratio of references to irrelevant information and the lowest ratio of references to the logical premises observed in the whole experiment.

EXPERIMENT 2

Experiment 2 was designed to replicate and extend Experiment 1. There are several difficulties of interpretation of Experiment 1 that Experiment 2 was intended to resolve. First, the belief bias observed could be due to the embedding of the logical premises in a prose passage. This could decrease the subjects' attention to the logically critical premises. Hence, a group was included that was given the premises only. Second, it is possible that the instruction to give verbal justifications, especially with a within-subjects design, could affect responses. Hence, another group was added with no instructions to verbalize. Another group was used to investigate further the causes of the trends in the protocol analysis. This group provided "thinking-aloud" protocols, which Ericsson and Simon (1980) argue are more likely to give an accurate picture of the informa-

tion heeded than is the retrospective method used in Experiment 1.

Finally, the problem structures were modified to take account of the "figural bias" discussed by Johnson-Laird and Steedman (1978). They show that the order in which terms are arranged in the syllogisms can exert an influence on the choice of conclusion, irrespective of logical validity. In Experiment 1, all premise pairs were of the form A-B, C-B, but valid conclusions were always of the form C-A and invalid conclusions of the form A-C. According to Johnson-Laird and Steedman's findings (but not their model), there may be a bias to prefer C-A conclusions with these premise types. This may have led us to overestimate subjects' logical ability in Experiment 1. In Experiment 2, both valid and invalid problems were associated with both A-C and C-A conclusions.

Method

Materials. The syllogisms and prose passages used were the same as those of Experiment 1, except for the modification that permitted both valid and invalid syllogisms to have both A-C and C-A conclusions. This was produced by interchanging the quantifiers of each of the original problem premises and reversing their conclusion, for example: "No A are B. Some C are B. Therefore, some C are not A" (original valid syllogism). "Some A are B. No C are B. Therefore, some A are not C" (valid control syllogism).

Design. As in Experiment 1, all subjects received four problems consisting of all four problem types combined with all four passage contents. In this experiment, four subject groups were used: Group 1 received prose passages and was required to verbalize the explanation for the decision retrospectively (as in Experiment 1). Group 2 was required to verbalize in the same manner as Group 1, but subjects received only the logical premises and not the full prose passage. Group 3 received prose passages and was required to verbalize concurrently (i.e., to think aloud while attempting to solve the problem). Group 4 received prose passages but was not required to verbalize at all.

Each of these four groups was then subdivided into two further groups, one of which received only A-C conclusions for both valid and invalid problems, and the other of which received C-A conclusions only.

Subjects. Sixty-four undergraduates at Plymouth Polytechnic acted as paid volunteers. They had no previous experience of this task and were tested individually.

Procedure. The instructions and problems were presented in the same manner as Experiment 1.

The instructions were as follows: Group 1—instructions as for Experiment 1. Group 2—instructions as for Experiment 1, except that any reference to the prose passage was omitted. Group 3—instructions as for first paragraph of Experiment 1, continued as follows, "Whilst you are trying to solve each problem I would like you to try to 'think aloud' as much as you can. Please don't let this distract you from the task in hand, which is to obtain the correct solution to the problem. If at any time during the task, I do not think that you are speaking enough, I will simply prompt you to speak a little more. Please take your time and be sure that you have the right answer before stating it. Any questions?" Group 4—instructions as for Experiment 1, except that any request to verbalize was omitted.

As in Experiment 1, subjects' protocols were recorded on a tape recorder for later analysis. Protocols were scored using the same procedure as in Experiment 1.

Results and Discussion

The percentage frequency of subjects accepting

arguments as valid in Experiment 2 is shown in Table 4. As is apparent from Table 4a, the order of terms in the conclusion had no significant effect on responses, and further analyses were collapsed over this factor. Binomial tests on the combined data of the subjects in all groups (n = 64) yielded highly significant evidence of the three effects found in Experiment 1. That is, more believable than unbelievable conclusions were accepted (p < .001), more valid than invalid conclusions were accepted (p < .001), and the two factors interacted (p < .01). The interaction reflects the fact that the belief-bias effect is more marked for invalid than for valid syllogisms. The interpretation of this interaction will be deferred to the general discussion. Overall, subjects were correct 87% of the time when logic accorded with belief and 48% of the time when it did not.

The same trends were manifested in each of the four groups (see Table 4b). In order to test whether response patterns were affected by group, a set of four 2 by 4 chi-square tests were carried out to compare yes/no frequencies across the four groups for each problem type. None of these analyses yielded a significant result. Clearly, the belief-bias effect is not due to embedding the premises in a prose passage, nor is it affected by instructions to verbalize, even if in a concurrent manner.

The results of the protocol analyses are shown in Table 5, for the three groups from whom protocols were collected. Inspection of Table 5 suggests that the distri-

bution of classification frequencies is, in fact, very similar for all three groups. This was confirmed statistically by rank ordering the 16 cell frequencies for each group and assessing the similarity of the rank orderings by Kendall's coefficient of concordance (cf. Siegel, 1956). There was high and significant concordance for both mention of the premises (W = .792, p < .001) and mention of irrelevant information (W = .871, p < .001). Consequently, further analyses were performed on the combined data of all three groups.

These analyses revealed highly significant interactions between the answer given to the problem and the protocol classifications for the valid-unbelievable problems. Subjects accepting the valid conclusion against its believability made more references to the logical premises ($\chi^2 = 24.61$, p < .001) and fewer references to irrelevant information ($\chi^2 = 15.11$, p < .001).

The other problem for which interactions were observed was the invalid-unbelievable type. Only six subjects went against both logic and belief to accept such arguments as valid. However, all six referred to both logical premises, and none referred to irrelevant information. Fisher exact probability tests revealed a significant interaction with the majority "no" responders in each case (p < .005 and p < .05, respectively). The simplest interpretation of these findings is that these subjects ignored beliefs and reasoned from the premises, but they did so with faulty logic.

Table 4
The Frequencies (in Percent) of Subjects Accepting Conclusions in Each Condition of Experiment 2

Conclusion	Problem Type									
	(a) All Groups (n = 64)					(b) Individual Groups (n = 16 in Each Group)				
	Valid		Invalid			Group	Valid		Invalid	
B	U	B	U	B	U		B	U	B	U
C-A	84	68	59	13	1	81	63	63	18	
A-C	88	56	72	13	2	81	63	44	6	
Mean	86	62	66	13	3	87	50	75	13	
					4	94	67	63	13	

Note—1 = prose passage, retrospective verbalization; 2 = premises only, retrospective verbalization; 3 = prose passage, concurrent verbalization; 4 = prose passage, no verbalization. B = believable; U = unbelievable.

Table 5
Classification Frequencies for Protocols of Experiment 2 (n = 16 in Each Group) Broken Down by Response Given

Problem Type		Both Premises								Irrelevant Information							
		Group 1		Group 2		Group 3		Combined		Group 1		Group 2		Group 3		Combined	
		Y	N	Y	N	Y	N	Y	N	Y	N	Y	N	Y	N	Y	N
Valid-Believable	M	<i>10</i>	0	6	1	5	1	21	2	3	1	8	0	6	1	17	2
	NM	3	3	7	2	9	1	19	6	10	2	5	3	8	1	23	6
Valid-Unbelievable	M	9	0	<i>10</i>	1	5	0	24	1	1	5	1	4	2	5	4	14
	NM	1	6	1	4	3	8	5	18	9	1	10	1	6	3	25	5
Invalid-Believable	M	4	2	6	2	3	0	13	4	6	2	4	2	8	1	18	5
	NM	6	4	4	4	9	4	19	12	4	4	6	4	4	3	14	11
Invalid-Unbelievable	M	3	5	1	9	2	1	6	15	0	6	0	5	0	7	0	18
	NM	0	8	0	6	0	13	0	27	3	7	1	10	2	7	6	24

Note—M = mentioned; NM = not mentioned; Y = yes response; N = no response. Correct responses are italicized.

The interpretation of the protocol data will be taken up in the general discussion. There is one further problem concerning the interpretation of the decision frequencies that Experiment 3 was designed to deal with.

EXPERIMENT 3

In both Experiments 1 and 2, subjects accepted significantly more valid than invalid conclusions. This suggests that people have some ability to reason and overcome belief biases, at least for unbelievable conclusions. It is, however, possible that this apparent logicity is an artifact of a response bias different from that controlled in Experiment 2. In all the syllogisms used so far, the quantifier "some" always modified the same term in the premises as in the conclusion for valid problems (e.g., "No A are B. Some C are B. Therefore, some C are not A."), but for invalid problems, this was never the case (e.g., "No A are B. Some C are B. Therefore, some A are not C.").

It is therefore possible that some form of feature-matching bias is responsible for the main effect of validity. This problem can be overcome if syllogisms in Figure 3 rather than Figure 2 are employed. For example, the following is logically equivalent to the former of the two problems above: "No B are A. Some B are C. Therefore, some C are not A."

Since in all Figure 3 syllogisms the two terms used in the conclusion (A and C) appear in the predicates of the premises, the possible response bias described could not operate. Experiment 3, therefore, compared subjects' performance on Figure 2 and Figure 3 syllogisms. Prose passages were employed, and thinking-aloud protocols were recorded.

Although the instructions presented in Experiments 1 and 2 clearly indicated that subjects' inferences should be based on logical necessity, this is an unusual requirement for subjects to follow. One interpretation of the belief-bias effect is that subjects "fail to accept the logical task" (Henle, 1962). In order to counter this possibility, the instructions of Experiment 3 were reworded to increase emphasis on the concept of logical necessity.

Method

Design. All subjects received eight problems to solve, the four types used previously in both Figure 2 and 3 syllogisms. Four further scenarios were constructed to add to the four used in Experiments 1 and 2, and each subject received each of the eight scenarios, randomly matched to the eight types of problems. The conclusion ratings for these additional problem contents are shown in Table 1. Presentation order was also randomized. All subjects were instructed to "think aloud" while solving the problems.

Subjects. Thirty-two 1st-year psychology students of Plymouth Polytechnic participated in partial fulfillment of course credit requirements. All were tested individually.

Procedure. The relevant section of the modified instructions follows: "Your task is to decide whether or not a given conclusion follows logically from the information given, and this information only. You must assume that all the statements within the passage are true; this is very important. If, and only if, you judge that the given conclusion logically follows from the statements given in the passage you should answer 'yes,' otherwise 'no.'"

The final sentence of this extract was repeated at the very end of the instructions. The procedure was otherwise similar to that of Experiments 1 and 2, with "thinking-aloud" instructions. Protocols were tape-recorded and subsequently transcribed and analyzed in a manner similar to that of the previous experiments.

Results

The frequencies of responses to the problems are shown in Table 6. It is apparent that there is no difference in performance between the Figure 2 and Figure 3 problems, thus eliminating the response-bias explanation of the validity effect. On the combined data, there were highly significant preferences to accept conclusions that were believable rather than unbelievable

Table 6
The Percentage Frequencies of Subjects Accepting Conclusions in Experiment 3

	Problem Type			
	Valid		Invalid	
	B	U	B	U
Figure 2	91	53	69	3
Figure3	91	53	66	9
Combined	91	53	67	6

Note—B = believable; U = unbelievable.

Table 7
Classification Frequencies for Protocols of Experiment 3 (n = 32) Broken Down by the Response Given

Problem Type		Both Premises				Irrelevant Information			
		Figure 2		Figure 3		Figure 2		Figure 3	
		Y	N	Y	N	Y	N	Y	N
Valid-Believable	M	5	1	11	0	16	1	17	1
	NM	24	2	18	3	13	2	12	2
Valid-Unbelievable	M	12	7	10	4	8	7	8	9
	NM	5	8	7	11	9	8	9	6
Invalid-Believable	M	7	6	3	4	12	5	14	5
	NM	15	4	18	7	10	5	7	6
Invalid-Unbelievable	M	1	8	1	7	1	20	2	16
	NM	0	23	2	22	0	11	1	13

Note—M = mentioned; NM = not mentioned; Y = yes response; N = no response. Correct responses are italicized.

($p < .001$, binomial test) and those that were valid rather than invalid ($p < .001$, binomial test). The interaction, although in the same direction as observed previously, fell just short of significance ($p = .067$, one-tailed binomial test). Subjects were correct 97% of the time when belief agreed with logic and 43% of the time when belief conflicted with logic.

The protocol analyses are summarized in Table 7. Previous experiments showed an interaction between the classification frequencies and type of response for the valid-unbelievable condition. A similar trend was observed in Experiment 3 on the references to premises criterion, although it fell short of significance for both figures. No interaction trend was apparent on the reference to irrelevant information criterion. No other conditions produced significant interactions.

GENERAL DISCUSSION

Over the three experiments, consistently large and significant effects of belief bias have been observed, despite controls for conversion of premises (cf. Revlin et al., 1980). Similarly, large and consistent effects of logical validity have been observed despite the controls introduced to test response-bias explanations in Experiments 2 and 3. There is also a consistent trend for the two factors to interact, such that the belief-bias effect is more marked for invalid than for valid problems. The strong instructional emphasis on logical necessity in Experiment 3 renders implausible any suggestion that the belief-bias effect reflects uncertainty on the subjects' part of what they were required to do. If they are "failing to accept the logical task," it is because they are unable to do so.

The Belief by Logic interaction arises because subjects respond differently to the two conditions in which logic and belief conflict. When the problem is invalid but believable, subjects generally accept the conclusion. Response rates are intermediate, however, when the syllogisms are valid but have unbelievable conclusions. This condition is especially interesting also with reference to protocol analysis, since subjects conforming to logic tend to refer to the premises, whereas those conforming to beliefs tend to refer instead to irrelevant information.

We must now ask what process of reasoning could account for these findings. There are a number of published models of syllogistic reasoning (e.g., Dickstein, 1978a, 1978b; Erickson, 1974; Guyote & Sternberg, 1981; Johnson-Laird & Steedman, 1978; Revlin, 1975a, 1975b). These models differ considerably in the details of their psychological descriptions, but in one respect they all agree. All the models suppose that the subject starts by forming a representation of the premises and then generates a conclusion, or set of possible conclusions, by a more or less logical (according to the model) process of reasoning. The subject then selects from the available list of conclusions one that matches the one

that he or she has generated (or, in the case of Johnson-Laird & Steedman, 1978, he or she simply writes down the conclusion generated).

None of these models can, in its present form, account for the results of the present study. The main sources of error permitted by these models are either faulty representation of the premises (conversion) or figural bias in the processing of representations. Our syllogisms were constructed such that all premises were legally convertible; the figures chosen were those least susceptible to figural bias and, in any case, were consistent across conditions. Finally, the results cannot be explained by atmosphere bias, either (cf. Begg & Denny, 1969; Woodworth & Sells, 1935), since all conclusions were equally favored by atmosphere.

Two of the models provide additional scope for the occurrence of reasoning errors; those of Guyote and Sternberg (1981) and Johnson-Laird and Steedman (1978). Only Guyote and Sternberg have made an attempt to explain the effects of problem content on reasoning. One of their experiments compared reasoning with factual (believable) and counterfactual (unbelievable) content, but they do not discuss possible interactions with validity. They do say that content affected reasoning and that the subsequent parameter estimations for their model suggest that "subjects store and manipulate factual information with greater ease than they do other kinds of information (Guyote & Sternberg, 1981, p. 499). This implies that subjects should reason more logically with believable than with unbelievable content. In fact, the interaction observed in the present experiments was the opposite of this: Subjects were more sensitive to logical validity on unbelievable problems.

We do not believe that our results can be explained on the assumption that all reasoning proceeds from the representation of the premises toward a conclusion. It appears that subjects not only check the validity of the conclusion (by reference to the premises) but are also influenced by a separate, direct assessment of its truth value. There are several ways in which the Belief by Validity interaction could arise. First, it may be that subjects accept uncritically a conclusion with which they agree but are more likely to check the logic if they do not agree with the conclusion. This is directly analogous to the finding of Lord, Lepper, and Ross (1979) that people will accept at face value the evidence of research studies whose conclusions agree with their prior beliefs, but they will criticize the design and methodology of those with conflict.

There are however, other explanations of the interaction. In the valid-unbelievable condition, the conflict is that logic dictates that the conclusion must be accepted despite its unbelievability. However, in the invalid-believable condition, the conflicting role of logic is less strong. Logically, the invalid conclusions do not necessarily follow from the premises, but neither are they contradicted by them. Since the conclusion is not

inconsistent with the premises, subjects may feel justified in favoring belief. Dickstein (1980, 1981) has presented evidence that subjects may indeed have difficulty in understanding this aspect of the concept of logical necessity. It is possible that the somewhat weaker Belief by Logic interaction observed in Experiment 3 was due to the modification of the instructions that emphasized that “yes” answers should be given if and only if the conclusions followed logically from the premises. This second explanation differs from the first in assuming that subjects always evaluate both the validity and the believability of each conclusion, but they respond differently to the two types of conflict.

A third explanation is that the conflict arising from unbelievable conclusions is less strong than that arising from believable conclusions, with our particular materials. Inspection of Table 1 reveals that while “true” statements are rated very close to the top of the scale, “false” items are rated, on average, only 1 point lower than the midpoint of the scale. Thus, in the valid-unbelievable condition, the bias to reject the conclusion on the basis of belief may be less strong than is the bias to accept, in accordance with belief, in the invalid-believable condition.

Unlike previous studies, we also have verbal protocol data to consider. The combined data of the three experiments are shown in two different ways in Table 8. Table 8a shows the probability of giving a particular explanation as a function of the response made. This is the appropriate way to look at the data if one assumes that they are rationalizations. In addition to the interaction that has been noted for the valid-unbelievable condition, a trend in the other conflict emerges on these pooled data. It seems that there is a tendency to give more references to irrelevant information when accepting invalid but believable conclusions. This suggests that subjects do perceive a conflict between logic and belief in this condition also, although less markedly than for valid-unbelievable problems.

On the other hand, if one supposes that the protocols do reflect the basis on which subjects were reasoning, then it is more appropriate to look at the likelihood of responses, given the protocol scores (see Table 8b). An interesting picture emerges here. It seems that the

Logic by Belief interaction is present for problems in which subjects refer to the premises or do not make irrelevant references, but it is absent when the premises are not cited or when irrelevant information is mentioned. The latter problems show almost pure belief bias, with little effect at all of validity.

This might seem to suggest that there are two kinds of subjects: some resting their conclusions on the premises, and others not. However, although there is generally a negative relation between scores on the two criteria on any particular subject’s response to a given problem, most subjects score positively on both criteria somewhere on their problems. It is still of interest to know whether subjects respond in a consistent manner to the conflict created by the valid-unbelievable condition. This can be examined by comparing the subjects’ responses to the Figure 2 and 3 problems of this type in Experiment 3. It turns out that the response rates on the two tasks are quite independent: Of the 16 accepting the Figure 2 conclusion, 9 accepted a Figure 3; of the 16 who did not accept a Figure 2, 8 accepted a Figure 3.

The above analysis supports the idea of a within-subjects conflict, as opposed to individual differences in strategies, which accords with Evans’ (1977b) discussion of the Wason selection task. Indeed, the whole pattern of results is consistent with the theory of reasoning put forward by Evans (1982) and previously applied to conditional reasoning problems. That is, response probabilities reflect competing logical and nonlogical processes. However, we still have two areas of uncertainty in the interpretation of the results: (1) Which explanation of the Belief by Logic interaction is correct? and (2) to what extent do the protocols reflect rationalizations, and to what extent the actual basis of subjects reasoning? An additional treatment of the protocols, shown in Table 9, helps to provide answers to both these questions.

This analysis was concerned with the order of mention of the premises and conclusion. The analysis was confined to “thinking-aloud” protocols only (Experiment 2, Group 3, and Experiment 3 combined), since it was thought that order of items in retrospective protocols need not reflect the actual order in which the subjects did things. There were three main classifica-

Table 8
Percentage of Positive Protocol Scores (Mentions) as a Function of Response Given and Percentage of Yes Responses as a Function of Protocol Scores

	(a) Percentage of Positive Protocol Scores				(b) Percentage of Yes Responses			
	Both Premises		Irrelevant Information		Both Premises		Irrelevant Information	
	Yes	No	Yes	No	M	NM	M	NM
Valid-Believable	41	31	47	36	91	86	91	86
Valid-Unbelievable	68	19	26	63	79	31	31	68
Invalid-Believable	30	41	60	43	70	77	81	69
Invalid-Unbelievable	72	23	45	54	25	3	5	11

Note—M = mentioned; NM = not mentioned. Data are combined for the three experiments; average Figure 2 and Figure 3 responses were used for Experiment 3 (n = 104).

Table 9
 Analysis of Thinking-Aloud Protocols for Experiment 2, Group 3 (n = 16), and for
 Experiment 3, Measured over Figures 2 and 3 (n = 32) Combined

	a				b		
	Conclusion Only	Conclusion to Premises	Premises to Conclusion	Other	Conclusion Only	Conclusion to Premises	Premises to Conclusion
Valid-Believable	39	29	25	7	86*	100*	88*
Valid-Unbelievable	34	28	34	3	70	48	30
Invalid-Believable	46	24	27	3	73	78	54
Invalid-Unbelievable	41	33	17	8	98*	97*	81*

Note—*a* = percentage frequency of protocol classifications as a function of problem type; *b* = percentage of decisions favoring belief on each problem as a function of protocol classification. *Logically correct response.

tions: Conclusions only (C)—These protocols refer to the conclusion but do not mention either premise. They may or may not include references to irrelevant information. Conclusion to premises (CP)—These protocols included reference to at least one premise after mention of the conclusion. Premises to conclusions (PC)—These protocols included mention of at least one premise followed by mention of the conclusion. (In some cases, CP and PC protocols also included irrelevant information.)

Table 9a shows that the great majority of protocols were classifiable in one of these three ways and that the distribution of classifications over the four problem types was quite similar. Table 9b shows the percentage of subjects favoring belief (saying “yes” on believable or “no” on unbelievable problems) as a function of protocol classification. Several features of this table suggest that the protocol analysis is diagnostic of the amount of logical reasoning subjects are doing. The PC protocols are associated with least belief bias in the two conflict conditions, suggesting that they do reflect more logical (premise-to-conclusion) reasoning. Note also that almost all logical errors in the invalid-unbelievable condition are associated with PC protocols. This confirms the interpretation given in discussion of Experiment 2 that such errors arise from subjects who ignore belief, reason from the premises, but make a logical error in doing so.

The C protocols, we suggest, are those of subjects who focus their attention on the conclusion and thus give the highest rates of belief bias. Even here, however, there are two sources of evidence that the premises, although not mentioned, have some influence on the subjects. One is the visible fact in Table 9b that their rate of favoring belief is higher if the response is also logically correct. The second is the fact that the mention of irrelevant information by such subjects (as inferred from Table 8) is greater when the logic contradicts the belief-biased response. Thus C protocols are primarily associated with belief-biased, rationalizing subjects.

The CP protocols reflect subjects who focus on the conclusion but also go on to consider the premises. They show intermediate rates of belief bias, presumably because in some cases the premises are seen to contra-

dict the belief, but they still do less well than those who attempt to reason from premises to conclusion.

This interpretation of Table 9b suggests that there is no singular answer to the problem posed by Table 8: The protocols partially reflect the basis of subjects' reasoning and partially reflect rationalization. We now ask what Table 9 can contribute to our understanding of the Belief by Logic interaction. The frequency of PC protocols is unaffected by the believability of the conclusion in Table 9a (26% with believable and 26% with unbelievable conclusions). This is not surprising if the PC protocols indicate subjects engaged in premise-to-conclusion reasoning. However, our first interpretation of the interaction does predict a shift between C and CP protocols for subjects who focus initially on the conclusion. Specifically, it was suggested that such subjects are more likely to go on to consider the premises if the conclusion is unbelievable than if it is believable. Some shift in this direction is actually observed. There were 43% C protocols for believable problems, dropping to 38% on unbelievables; there were 27% CP protocols on believables, rising to 31% on unbelievables.

The general picture of Table 9 does not, however, support this interpretation of the interaction, for two reasons. First, the observed shift between C and CP protocols is too small to account for the large interaction in the response frequencies. Second, both CP and PC protocols are associated with substantially more belief bias in the invalid-believable condition than in the valid-unbelievable condition. This must mean that subjects who take account of both logic and belief experience more competition from belief in the former condition than in the latter. Thus, either our second or our third interpretation of the interaction is to be preferred. Whether it is due to weaker logic on invalid than on valid problems or to stronger belief on believable than on unbelievable problems cannot be distinguished in the present experiments.

Finally, the analysis shown in Table 9 is helpful in resolving a problem raised by a reviewer of this paper, namely, that believability of the conclusion is inevitably confounded with believability of the premises. By the laws of logic, any valid argument with a false conclusion

must have either a false premise or premises with incompatible suppositions. Such premises might strike the subject as anomalous and thus inhibit reasoning. However, the data of Table 9 strongly suggest that this is not the main cause of belief bias on valid-unbelievable problems. Table 9a shows that most subjects use a conclusion-centered strategy on these problems, and Table 9b shows that the majority (70%) of those who do reason from premises to conclusion correctly accept the inference.

In conclusion, we hope that we have shown that the introduction of protocol analysis has proved to be a most productive way of differentiating and understanding the processes underlying the belief-bias effect. The picture that finally emerges is that logic and belief conflict throughout, but they do so at different levels. When subjects focus primarily on the conclusion, belief biases are maximal and logical effects are minimal. Such subjects often rationalize their responses by referring to irrelevant information. These are the clearest examples of Henle's (1962) "failure to accept the logical task," but they still show a small influence of logic. On about 25% of occasions, though, a genuine premise-to-conclusion inference is attempted, with much higher logical success. It is most important to note, however, that even in cases in which the logical task is accepted, substantial (although lesser) effects of belief bias are still observed. These findings not only provide a challenge for existing models of syllogistic reasoning but also raise broader questions about people's rational competence to generate and assess logical arguments in real life, whenever they have clear a priori beliefs about the subject under discussion.

REFERENCES

- BEGG, I., & DENNY, J. P. Empirical reconciliation of atmosphere and conversion interpretations of syllogistic reasoning errors. *Journal of Experimental Psychology*, 1969, **81**, 351-354.
- BRAINE, M. D. S. On the relation between the natural logic of reasoning and standard logic. *Psychological Review*, 1978, **85**, 1-21.
- COHEN, L. J. Can human irrationality be experimentally demonstrated? *Behavioural and Brain Sciences*, 1981, **4**, 317-370.
- DICKSTEIN, L. S. The effect of figure on syllogistic reasoning. *Memory & Cognition*, 1978, **6**, 76-83. (a)
- DICKSTEIN, L. S. Error processes in syllogistic reasoning. *Memory & Cognition*, 1978, **6**, 537-543. (b)
- DICKSTEIN, L. S. Inference errors in deductive reasoning. *Bulletin of the Psychonomic Society*, 1980, **6**, 414-416.
- DICKSTEIN, L. S. Conversion and possibility in syllogistic reasoning. *Bulletin of the Psychonomic Society*, 1981, **18**, 229-232.
- ERICKSON, J. A set analysis theory of behavior in formal syllogistic reasoning tasks. In R. L. Solso (Ed.), *Theories of cognitive psychology: The Loyola Symposium*. Hillsdale, N.J.: Erlbaum, 1974.
- ERICSSON, K. A., & SIMON, H. A. Verbal reports as data. *Psychological Review*, 1980, **87**, 215-251.
- EVANS, J. St. B. T. Linguistic factors in reasoning. *Quarterly Journal of Experimental Psychology*, 1977, **29**, 297-306. (a)
- EVANS, J. St. B. T. Toward a statistical theory of reasoning. *Quarterly Journal of Experimental Psychology*, 1977, **29**, 621-635. (b)
- EVANS, J. St. B. T. *The psychology of deductive reasoning*. London: Routledge & Kegan Paul, 1982.
- FALMAGNE, R. J. (Ed.). *Reasoning: Representation and process*. Hillsdale, N.J.: Erlbaum, 1975.
- FEATHER, N. T. Acceptance and rejection of arguments in relation to attitude strength, critical ability and intolerance of inconsistency. *Journal of Abnormal and Social Psychology*, 1964, **69**, 127-36.
- GORDEN, R. Attitudes toward Russia on logical reasoning. *Journal of Social Psychology*, 1953, **37**, 103-111.
- GUYOTE, M. J., & STERNBERG, R. J. A transitive-chain theory of syllogistic reasoning. *Cognitive Psychology*, 1981, **13**, 461-525.
- HENLE, M. On the relation between logic and thinking. *Psychological Review*, 1962, **69**, 366-378.
- HENLE, M., & MICHAEL, M. The influence of attitudes on syllogistic reasoning. *Journal of Social Psychology*, 1956, **44**, 115-127.
- INHOLDER, B., & PIAGET, J. *The growth of logical thinking*. New York: Basic Books, 1958.
- JANIS, I., & FRICK, P. The relationship between attitudes toward conclusions and errors in judging logical validity of syllogisms. *Journal of Experimental Psychology*, 1943, **33**, 73-77.
- JANIS, I., & TERWILLIGER, R. An experimental study of psychological resistances to fear arousing communications. *Journal of Abnormal and Social Psychology*, 1962, **65**, 403-410.
- JOHNSON-LAIRD, P., & STEEDMAN, M. The psychology of syllogisms. *Cognitive Psychology*, 1978, **10**, 64-99.
- KAUFMAN, H., & GOLDSTEIN, S. The effects of emotional value of conclusions upon distortion in syllogistic reasoning. *Psychonomic Science*, 1967, **7**, 367-368.
- LEFFORD, A. The influence of emotional subject matter on logical reasoning. *Journal of General Psychology*, 1946, **34**, 127-151.
- LORD, C., LEPPER, M. R., & ROSS, L. Biased assimilation and attitude polarization: The effect of prior theories on subsequently considered evidence. *Journal of Personality and Social Psychology*, 1979, **37**.
- MORGAN, J., & MORTON, J. The distortion of syllogistic reasoning produced by personal convictions. *Journal of Social Psychology*, 1944, **20**, 39-59.
- NISBETT, R. E., & ROSS, L. *Human inference: Strategies and short-comings of social judgment*. Englewood Cliffs, N.J.: Prentice-Hall, 1980.
- NISBETT, R. E., & WILSON, T. D. Telling more than we can know: Verbal reports on mental processes. *Psychological Review*, 1977, **84**, 231-259.
- POLLARD, P. *Human reasoning: Logical and nonlogical explanations*. Unpublished doctoral thesis, Plymouth Polytechnic, 1979.
- POLLARD, P. Human reasoning: Some possible effects of availability. *Cognition*, 1982, **12**, 65-96.
- REVLIN, R. Syllogistic reasoning: Logical decisions from a complex data base. In R. Falmagne (Ed.), *Reasoning: Representation and process*. Hillsdale, N.J.: Erlbaum, 1975. (a)
- REVLIN, R. Two models of syllogistic reasoning: Feature selection and conversion. *Journal of Verbal Learning and Verbal Behavior*, 1975, **14**, 180-195. (b)
- REVLIN, R., & LEIRER, V. The effects of personal biases on syllogistic reasoning: Rational decisions from personalized representations. In R. Revlin & R. E. Mayer (Eds.), *Human reasoning*. Washington, D.C.: Winston-Wiley, 1978.
- REVLIN, R., LEIRER, V., YOPP, H., & YOPP, R. The belief-bias effect in formal reasoning: The influence of knowledge on logic. *Memory & Cognition*, 1980, **8**, 584-592.
- REVLIN, R., & MAYER, R. E. (Eds.). *Human reasoning*. Washington, D.C.: Winston-Wiley, 1978.
- SIEGEL, S. *Nonparametric statistics*. New York: McGraw-Hill, 1956.

- WASON, P. C., & EVANS, J. ST. B. T. Dual processes in reasoning? *Cognition*, 1975, 3, 141-154.
- WILKINS, M. The effect of changed material on ability to do formal syllogistic reasoning. *Archives of Psychology*, 1928, 16, 83.
- WILSON, W. The effect of competition on the speed and accuracy of syllogistic reasoning. *Journal of Social Psychology*, 1965, 65, 27-32.
- WINTHROP, H. Semantic factors in the measurement of personality integration. *Journal of Social Psychology*, 1946, 24, 149-175.
- WOODWORTH, R. S., & SELLS, S. B. An atmosphere effect in syllogistic reasoning. *Journal of Experimental Psychology*, 1935, 18, 451-460.

(Received for publication July 6, 1982;
revision accepted January 31, 1983.)