Prevarication: Reasoning from false assumptions*

RUSSELL REVLIS†

California State University, Fullerton, California 92634

This paper examines the strategies used in solving prevaricative reasoning problems. The task consisted of a noncontradictory set of propositions followed by an assumption that introduced an inconsistency. Ss were required to restore consistency by assigning truth values to propositions. When reasoning from false assumptions, Ss exhibit a specific strategy in assigning truth values: If forced to choose between the truth of a generality (e.g., All As are Bs) and that of a particular fact (e.g., This Y is a Z), the Ss assign the generality TRUE and the fact FALSE. Affirmative generalities are assigned TRUE more often than are negative ones, and those expressing class inclusion are assigned TRUE more often than those expressing property assignment. The data show that the selection of a "path" through this hypothetical reasoning problem cannot be attributed to either the syntactic form or the preexperimental credibility of the propositions. A generality coding model was discussed in terms of strategies shown in prevaricative reasoning problems as well as in other inference tasks.

Prevarication is a form of hypothetical reasoning where we *assume*, for the sake of argument, that a proposition is TRUE when, in fact, we *know* it is FALSE. Such is the case, for example, in using reductio ad absurdum arguments in mathematics where the truth of a proposition is tested by assuming its negation is true and then deriving some logical contradiction from the assumption. This form of conjecture is found in a number of scientific endeavors, as in contingency planning (Jouvenal, 1967), causal analysis, and theory testing (Simon, 1952, 1965; Simon & Rescher, 1966). Current studies of prevarication have demonstrated the operation of specific inference processes in conjectural reasoning (Revlis & Hayes, 1972; Revlis, Lipkin, & Hayes, 1971): When people reason from false assumptions, they exhibit consistent strategies in the assignment of truth values to propositions. These strategies appear to be independent of the superficial form of the propositions and suggest that the reasoner is attending to deeper aspects of the problems. Specifically, Revlis and Hayes (1972) showed that the truth-value assignments reflect the relative ability of propositions to *structure* the problem's elements, that is, their ability to reduce the uncertainty of the location of objects in classes (cf. Bartlett, 1932). The present experiments test a less tenable claim, that the truth-value assignments in this hypothetical reasoning task may not be the result of inferential processes, but rather they reflect response biases due to the syntactic form and the preexperimental credibility of the problem's statements.

The present task is intended as paradigmatic of the kinds of decision problems that individuals face when they engage in hypothetical reasoning. The task environment consists of (a) a noncontradictory set of propositions, representing the reasoner's knowledge, and

(b) a counterfactual assumption, which introduces inconsistencies that the S must resolve by reevaluating his knowledge (i.e., by assigning truth values to the statements). For example:

- Propositions: (1) This mouse is not a snake.
(2) All snakes are reptiles.
(3) This mouse is not a reptile.
Assumption: (4) This mouse is a snake.

The problem solver is to accept the assumption and to assign a truth value of FALSE to any of the statements that lead to a logical contradiction when combined with the assumption. Clearly, Statement 1 must be FALSE because it directly contradicts the assumption. At this point, the reasoner "knows" that in this hypothetical domain "This mouse is a snake," but it is unclear whether the mouse is a mammal or a reptile. To resolve this problem, the reasoner may choose from two equally logical paths. In the first, which may be called the *generalist* path, the generality (Statement 2) is assigned TRUE, which demands that the particular fact (Statement 3) be assigned FALSE. This last truth-value assignment is required because the S knows that this mouse is a snake (the assumption) and that all snakes are reptiles (Statement 2); therefore, Statement 3 must be FALSE when it claims that this mouse is not a reptile. In the second solution, called the *particularist* path, the particular fact (Statement 3) is accepted as TRUE, which demands that the generality (Statement 2) be FALSE. This is because the S knows that this mouse is a snake (the assumption) and that this mouse is not a reptile (Statement 3); therefore, Statement 2 is FALSE when it claims that all snakes are reptiles.

Despite the fact that either solution is equally logical, people express a preference for the *generalist* solution. Revlis et al (1971) showed that, when forced to choose between the truth of a generality (e.g., *All As are Bs*) and that of a particular fact (e.g., *This Y is a Z*), Ss tend to accept the generality and reject the fact. This

*This research was supported by a National Institute of Mental Health postdoctoral fellowship (1-F02-MH50536-01).

†The author wishes to express his appreciation to P. A. Carpenter, H. H. Clark, M. A. Just, N. J. Revlis, and E. E. Smith for their help in the preparation of this manuscript.

preference is more pronounced for generalities that affirm a relation than those that deny one: Affirmative generalities are assigned TRUE more often than are negative ones. Those generalities expressing a class-inclusion relation (e.g., *All As are Bs*) are assigned TRUE more often than those expressing property assignment (e.g., *All As have Cs*).

These data were predicted by a generality coding model based on a modal logic analysis of the propositions (cf. Rescher, 1964; Revlis & Hayes, 1972). The model claims that when Ss are asked to conjecture about new relationships between objects they assign TRUE to those propositions that provide an ordering rule within the domain of the problem (i.e., the most general predicates). This is accomplished in two phases. First, the reasoner is said to construct a possible world in which the assumption can be TRUE. Second, the reasoner orders the available descriptions about the world in terms of their ability to structure to domain, that is, in terms of modal categories (Rescher, 1961, 1964). The more general the propositions (that is, the more instances it predicts in a well defined universe of discourse), the lower the modal value it has, the more central the statement will be for reasoning about new possible worlds.

Although this model has been shown to be a good predictor of the reasoner's decisions on these problems, it is possible that the decisions may be a result of two kinds of biases: first, a bias in the construction of the problems, such that the generalities were more consonant with the S's knowledge than were the particular facts. In this case, the truth values might really reflect credibility judgments and not the kind of inference processes that the generality coding model claims.

Second, one or more response biases may operate. There are at least two such biases pertinent to the present task: (1) a preference for affirmative statements and (2) a preference for general statements. Such biases (or preferences) may be independent of the kind of inferences that the reasoning task requires. That is, if Ss were simply asked to assign the words TRUE and FALSE to statements, without regard to the inductive or deductive consequences, the data might correspond to the previous findings for the hypothetical reasoning task.

To deal with these possible interpretations, the present experiments consider the importance of three factors for the selection of a path through this hypothetical reasoning problem: (1) the problem solver's belief in the credibility of the propositions, (2) the polarity (affirmative or negative) of the statements, and (3) the instructions to "reason" on this task. With respect to the first factor, Revlis et al (1971) tested the importance of the propositions' credibility by comparing the reasoner's decisions on *familiar* problems that contained real subjects and predicates (e.g., *All snakes are reptiles*) with decisions on *unfamiliar* problems, where the subjects and predicates were nonsense words

(e.g., *All trups are dorbes*)—a method used by Wilkins (1928). Such statements are at best amphiborous (cf. Steinberg, 1971) so that no meaningful credibility value can be assigned to them. Two results of the Revlis et al study are of interest here. First, the Ss assigned generalities with nonsense subject and predicate terms TRUE more often than they did the contrasting particular facts. Second, the Ss assigned TRUE to familiar generalities more frequently than they did to unfamiliar ones. The first result shows that the Ss prefer to reason with generalities whether or not the statements have any meaningful credibility. However, the second result suggests the possibility that, when the proposition's credibility is meaningful, its credibility does enter into the S's decision process. For example, the Ss may have preferred to assign familiar generalities TRUE because the problems were so constructed that the generalities had greater credibility than the particular facts. This alternative is tested in Experiment I by having the Ss make credibility judgments about the problem statements and then comparing these judgments to the Ss' truth-value assignments.

Revlis and Hayes (1972) showed that the assignment of truth values to generalities is independent of the polarity of the contrasted particular facts. That is, affirmative generalities (e.g., *All trees on the square are elms*) are assigned TRUE equally as often when they are contrasted with affirmative facts (e.g., *This tree is a pine*) as with negative facts (e.g., *This tree is not an elm*). Therefore, the Ss could not be simply making their truth-value assignments based on the polarity of the statements. However, the results of that study are not conclusive because (a) negative generalities were not tested against both affirmative and negative facts and (b) the context of the problems—the main independent variable—reduced the frequency of assigning generalities TRUE, which may have masked the importance of the polarity of the particular fact due to a regression effect. A more thorough examination of the importance of polarity of the propositions is presented in Experiment I.

The response bias toward accepting generalities is examined in Experiment II by employing a selection task in which the Ss are instructed to assign truth values independently of the inferential consequences—a procedure employed by Van Hook, Rose, & Taplin.¹

EXPERIMENT I

Method

Design and Procedure

The experiment consisted of two parts. In Part 1 the Ss were required to solve hypothetical reasoning problems of the kind illustrated above. In Part 2 the Ss assigned credibility values to each of the propositions found in the reasoning problems. In Part 1 each S received a booklet of 20 problems. The E read the instructions aloud while the Ss followed with their own copies of the instructions. A sample problem was presented and the two

possible solutions were demonstrated. The Ss were told to proceed through the problems at their own rates and to work on each problem in the order presented in the booklet (one problem per page). When Ss questioned the logical ambiguity of the solutions, they were instructed to choose the "best" alternative.

The generalities were either affirmative (e.g., All snakes are reptiles) or negative (e.g., No snakes are mammals). Similarly, the particular facts were either affirmative (e.g., This mouse is a mammal) or negative (e.g., This mouse is not a reptile). The two levels of polarity for the generality and for the fact were crossed to form four problem types. Each S saw five instances of each problem type. The instances were constructed from definitions given for object and animal names in the *Random House Dictionary of the English Language* (1966).

There were two groups of 26 Ss, corresponding to two levels of relation: class inclusion and property assignment. Class-inclusion relation was expressed with the plural form of the copula *is a* (e.g., All snakes are reptiles); property assignment was expressed with the plural form of *has a* (e.g., All birds have wings). Within a group, all Ss received the same sequence of problems. The sequence was formed by dividing the problems into five blocks, each of which was filled randomly with an instance of each problem type. On every problem, the assumption was positioned as the last statement in the problem. The order of the remaining three statements was randomized for each problem.

In Part 2 the Ss evaluated the credibility of each of the propositions. After an S solved the problems in his booklet, he was asked to complete a questionnaire, which consisted of all of the propositions found in his booklet. The Ss were to indicate whether they believed that a statement was (a) always true, (b) usually true, (c) usually false, (d) always false, or (e) uncertain. There were two questionnaires corresponding to the two kinds of relation. An additional two groups of 13 Ss completed a questionnaire without having seen any of the problems. This procedure permitted a determination of whether working the problems influenced the credibility judgments.

Subjects

The Ss were 78 Stanford University students fulfilling a course requirement for introductory psychology. The Ss were run in groups of approximately eight, in sessions lasting about 40 min.

Results

The data from the reasoning task are summarized in Tables 1 and 2 in terms of the percentage of problems on which generalities are assigned TRUE. If Ss make their decisions randomly, generalities should be assigned TRUE on only 50% of the problems. The results show that Ss exhibit a pronounced preference for reasoning with general propositions. Over all conditions, generalities were assigned TRUE significantly more often

Table 2
Percentage Acceptance of Universally Quantified
Property-Assignment Propositions*

Generality	Particular Facts		Total
	Affirmative	Negative	
Affirmative†	64.4	75.4	70.0
Negative††	52.5	63.5	58.0
Total	58.5	69.5	64.0

*Each cell represents the average of five decisions for 26 Ss.

†All cells associated with affirmative generalities show an acceptance rate significantly greater than chance ($p < .01$).

††Negative generalities contrasted with affirmative facts ($n.s.$); negative generalities contrasted with negative facts ($p < .05$).

than chance (Wilcoxon matched pairs, signed ranks: $p < .001$). This preference for generalities is larger when they express class-inclusion relation than property assignment (72.5% vs 64.0%, difference not significant). Affirmative generalities are assigned TRUE more often than negative generalities [$F(1,50) = 18.3$, $p < .001$]. This overall preference for affirmative generalities does not interact with the relation expressed by the propositions.

For class-inclusion problems, the assignment of truth values to generalities is independent of the polarity of the contrasting facts, thereby replicating the overall findings of Revlis and Hayes (1972). However, when reasoning with property-assignment relations, the truth values assigned are influenced by the polarity of the propositions. Although this group of Ss preferred to reason more with generalities than with facts, the generalities (both affirmative and negative) were assigned TRUE more often when they were contrasted with negative facts than with affirmative ones [$F(1,25) = 10.1$, $p < .01$].

The analysis of credibility values was performed on the questionnaire responses of 76 of the 78 Ss (two Ss failed to follow instructions). A comparison of the number of generalities and facts assigned to each of the five response categories showed that, for any credibility category, control Ss do not differ from those in the experimental group on their credibility judgments. Therefore, having performed the reasoning task does not contaminate the credibility judgments of the experimental groups.

The importance of credibility for the truth-value assignments was determined in three ways: (1) by comparing the credibility judgments for propositions that have different truth-value assignments; (2) by computing the correlation between credibility judgments and truth-value assignments over all problems²; (3) by determining the conditional probability of assigning a generality TRUE given specified credibility judgments for the contrasting statements within each problem. First, the data show that the differences found in truth-value assignments are not present in the credibility judgments. Tables 3 and 4 give the percentage of class-inclusion and property-assignment propositions assigned to each of the credibility categories by the

Table 1
Percentage Acceptance of Universally Quantified
Class-Inclusion Propositions*

Generality	Particular Facts		Total
	Affirmative	Negative	
Affirmative†	79.0	79.2	79.1
Negative	65.0	66.7	65.9
Total	72.0	73.0	72.5

*Each cell represents the average of five decisions for 26 Ss.

†All values are significantly greater than chance ($p < .01$).

Table 3
Percentage of Class-Inclusion Propositions in
Five Categories of Credibility

	Always True	Usually True	Undecided	Usually False	Always False
Generalities					
Affirmative	70.0	26.2	2.3	0.0	1.5
Negative	75.4	20.0	3.1	1.5	0.0
Total	72.7	23.1	2.7	0.8	0.8
Facts					
Affirmative	64.6	29.2	4.6	1.5	0.0
Negative	70.8	17.7	5.4	4.6	1.5
Total	67.7	32.4	5.0	3.0	0.8
Overall	70.2	23.2	3.8	1.9	0.8

Table 4
Percentage of Property-Assignment Propositions in
Five Categories of Credibility

	Always True	Usually True	Undecided	Usually False	Always False
Generalities					
Affirmative	79.2	12.3	4.6	1.5	2.3
Negative	89.2	4.6	3.1	1.5	1.5
Total	84.2	8.4	3.9	1.5	1.8
Facts					
Affirmative	75.4	10.8	10.0	1.5	2.3
Negative	73.8	8.5	6.2	1.5	10.0
Total	74.6	9.7	8.1	1.5	6.2
Overall	79.4	9.1	6.0	1.5	4.0

experimental groups. A comparison of the number of generalities and facts within each category shows that there are no differences between (a) affirmative propositions and negative ones and (b) generalities and facts. Overall, 75.5% of the generalities and 71.2% of the particular facts were judged to be highly credible (always true).

Second, the credibility judgments of generalities and facts within each problem are highly correlated (class inclusion: $r = .91$, $p < .001$; property assignment: $r = .82$, $p < .001$). In contrast, neither the credibility of the generalities nor the facts correlate significantly with the Ss' truth-value assignments for either relation (class inclusion: generality, $r = .14$, fact, $r = .05$; property assignment: generality, $r = .17$, fact, $r = .10$). Further, there is no significant multiple correlation between credibility *per se* and truth-value assignments (class inclusion: $R_{1.23} = .26$; property assignment: $R_{1.23} = .44$).

Third, although the data show that the priority accorded generalities is independent of the credibility judgments for the entire set of propositions, it is still possible that on any specific problem S's judgments are influenced by the differential credibility of the contrasting statements. To examine this possibility, Table 5 gives the conditional probabilities of assigning a generality TRUE, given the specified credibility values, averaged across problem types. Only those credibility

categories that contain 2% or more of the decisions for each condition are included. As a result, only two levels of credibility were tabulated (always true and usually true); they represent the dominant response categories (class inclusion: 94.0% of all decisions; property assignment: 88.2% of all decisions). Owing to the disproportionate cell frequencies, the values given in Table 5 were analyzed by the method of expected cell frequencies (Myers, 1966). The results show that for both relations, the probability of assigning a generality TRUE is independent of both its own credibility (always true generalities do not differ from usually true ones) and that of its contrasted particular fact (always true facts do not differ from usually true ones). However, for class-inclusion statements, there is a significant interaction of Acceptance Probability of the Fact by Generality [$F(1,154) = 73.5$, $p < .001$]. A generality is more likely to be assigned TRUE (for class-inclusion statements) when it is contrasted with a fact of lower credibility than itself. This increase in the number of generalities assigned TRUE in class-inclusion problems makes only a marginal contribution to the overall preference data: Lower credibility of facts contributes to acceptability on only 3.6% of all decisions on class-inclusion problems. Interestingly, the assignment of TRUE to facts is not increased when they are contrasted with generalities of lower credibility than themselves. When facts have a greater prior credibility than generalities (class inclusion: usually true generalities contrasted with always true facts), the generalities are still assigned TRUE more often than chance ($p < .001$).

For property-assignment problems, the conditional probability of assigning generalities TRUE is constant across credibility values. The preference for reasoning with generalities on these problems only exceeds chance when both the fact and the generality are believed to be always true ($p < .001$). The frequency of assigning generalities TRUE is clearly not enhanced when the contrasting facts have a lower prior credibility.

Discussion

The present findings indicate that, when Ss reason from false assumptions, they accord generalities a priority over particular facts. It is clear that the pattern

Table 5
Conditional Probability of Accepting Universally
Quantified Propositions*

Relation	Generality	Particular Facts	
		Always True	Usually True
Class	Always True	.74 (71.2)	.95 (3.6)
	Usually True	.73 (6.4)	.74 (12.8)
Property	Always True	.62 (76.4)	.63 (4.8)
	Usually True	.62 (2.6)	.56 (4.4)

*Numbers in parentheses indicate the percentage of problems where the propositions have the tabulated credibility values.

of truth-value assignments in the present experiment cannot be attributed to the differential credibility of the contrasting statements. The Ss' assignment of truth values is similar across conditions where statements have either no prior credibility (unfamiliar problems, Revlis et al, 1971) or have a range of credibility values, as in the present experiment. This is not to say that credibility will be unimportant when the statements within a problem are widely disparate with respect to credibility. The results suggest only that, for the problems in this and previous studies, the credibility of the propositions is not sufficient to account for the reasoners' decisions.

Although the polarity of the statements is shown to influence which truth values will be assigned to the propositions, Ss do not make their decisions solely on the basis of which statement is affirmative. The effect of polarity varies with the relation that the propositions express. For class-inclusion relations, S's judgments are independent of the polarity of the contrasting statements. However, for property-assignment relations, the priority accorded generalities changes as a function of the contrasting statement's polarity in two ways. First, when there is a polarity difference between the contrasting statements, the affirmative propositions are given an increased priority (approximately 11%, see Table 2). Second, when there is no polarity difference between statements, there is no difference in the frequency with which affirmative and negative generalities are assigned TRUE (both are assigned TRUE more often than chance).

The different patterns of decision making between these two relations pose particular problems for a global explanation of prevaricative reasoning. Why, for example, should property-assignment problems result in a greater number of particularist decisions than do class-inclusion problems? Such questions can be best answered when we have a model of the reasoner's representation of the propositions as well as the kinds of operations he performs on that representation. Before considering a first approximation to such an inference model, it is necessary to examine another form of the "response bias" explanation.

EXPERIMENT II

Method

Although the results of Experiment I indicate that the decisions on these problems cannot be attributed solely to the credibility of the statements or to a response bias toward affirmative statements, it still remains necessary to consider whether Ss are actually *reasoning* on these problems. To explore this, two groups of Ss were presented with the class-inclusion problems of Experiment I. Group I was given the reasoning instructions as described above, while Group II was given selection instructions. That is, they were told to assign the directly contradictory statement FALSE and to assign one of the remaining two statements FALSE and the other TRUE, at their own whim. The Ss were told that the E was interested in whether they could find the directly contradictory statement.

Table 6
Instructions to Reason and Percentage Acceptance of Generalities

Instruc- tions	Generality	Particular Facts		Total
		Affirm- ative	Neg- ative	
Reasoning	Affirmative	84.2	90.0	87.1
	Negative	65.0	71.7	68.4
	Total	74.6	80.9	77.8
Selection	Affirmative	68.3	83.3	75.9
	Negative	20.0	51.5	35.8
	Total	44.2	67.4	55.8

These selection instructions promote a modest degree of analysis of the propositions in the problems (finding direct contradictions minimally entails determining the logical subject and object and whether the relation is affirmative or negative). Yet, these instructions fall short of demanding inductive or deductive inference, which presumably distinguishes these instructions from the reasoning instructions.

The Ss were 48 California State University students fulfilling a course requirement for introductory psychology. The Ss were run in groups of approximately four, in sessions lasting about 30 min.

Results

Table 6 summarizes the performance of the reasoning and selection groups. It shows that there is no overall difference in performance between Group I (the reasoning group) and its corresponding group in Experiment I (class-inclusion problems). Despite the slightly higher frequencies of assigning generalities TRUE (difference nonsignificant), Group I can be characterized in the same manner as previous reasoning groups. That is, these Ss show an overall preference for reasoning with generalities. Although affirmative generalities are assigned TRUE more frequently than are negative generalities [$F(1,23) = 16.8$, $p < .001$], both are assigned TRUE more often than chance ($p < .001$). Further, the preference for affirmative generalities is independent of the polarity of the contrasting particular facts.

The characterization of the decisions of the selection group is in marked contrast to that of the reasoning group. First, the overall frequency of assigning generalities TRUE for the selection group is not significantly different from what would be expected by chance alone. Second, the polarity of the statements is critical for the decisions. Not only are affirmative generalities assigned TRUE more often than negative ones [$F(1,23) = 127.7$, $p < .001$] but also the assignment frequencies differ as a function of the polarity of the contrasting particular facts: Affirmative facts are assigned TRUE more often than are negative ones [$F(1,23) = 26.9$, $p < .001$]. In general, when making truth-value assignments under the selection instructions, if one of the statements is affirmative, it will be assigned TRUE more often than chance, independent of whether the statement is a generality or

a particular fact ($p < .001$). Indeed, when negative generalities are contrasted to affirmative particular facts, the negative generalities are assigned FALSE significantly more often than chance ($p < .01$). Clearly, when Ss are given full reasoning instructions, their decisions are different than when they are not instructed to reason.

Discussion

The results of Experiment II show that, when Ss are instructed to consider the inductive and deductive consequences of propositions, the pattern of decisions is different than when such complex inferences are not called for. It is plausible that the different patterns of decisions of the selection and reasoning groups are attributable to the fact that complex inferences were required of the reasoning group but not of the selection group. However, such claims cannot be made with full confidence. Certainly some analysis of the propositions was present in the selection group as well as in the reasoning group. This is attested by the strong preference for affirmative propositions in the former group; such preferences for propositions that affirm a relation could not be exhibited without some analysis of the subject and object classes and the relation that holds between them. Certainly such an analysis would be called for when Ss searched for the proposition that directly contradicted the assumption (all Ss were 100% successful in completing this portion of the task). Yet, it is equally clear that the selection and reasoning groups differ with respect to their overall performance [$F(1,46) = 7.33, p < .01$]. Therefore, while the selection group cannot be said to be devoid of reasoning processes on these problems, they cannot be claimed to be making the full inferences that the reasoning group was instructed to make.

One way to characterize the decisions of the selection group is to claim that their instructions induced a response bias for assigning affirmative statements TRUE and, to a small degree, a bias for affirmative generalities. However, to claim that similar biases operate under the reasoning instructions is both contrary to the data from the two experiments and unparsimonious: Such a response bias would have to entail a complex interaction between the relations expressed by the propositions, the polarity of that relation (affirmative or negative), the contrasting polarity of other propositions in the set, and, finally, the explicit quantification of each proposition in the set (cf. Revlis & Hayes, 1972 for an analysis of the attributes of such models).

GENERAL DISCUSSION

The present experiments lend support to the hypothesis that the Ss' decisions on prevaricative reasoning problems are the result of inferential processes and not the reflection of simple response biases. That is,

such decisions cannot be attributed to the reasoner's belief in the credibility of the propositions or to the polarity of the contrasting propositions. Support is also given to the hypothesis that Ss bring to bear inferential processes when solving the reasoning problems. Consequently, it is useful to consider how S's knowledge is reordered on these problems in order to bring consistency to the hypothetical domain described by the propositions. That is, how can we characterize the reasoner's knowledge and the operations he performs on that knowledge in order to account for his decisions?

A first approximation to a representation of the reasoner's knowledge in the problems is the generality coding model. This claims that the reasoner's real-world knowledge is hierarchically organized by predicates of varying degrees of generality, somewhat akin to Kintsch's model of semantic memory (Kintsch, 1972). The structure of this knowledge need not be explicitly specified as propositions but might be equivalently represented as a branching tree structure (Collins & Quillian, 1969, 1970a, b) or a set-theoretic representation (Meyer, 1970). When the S is asked to conjecture about new relationships between objects, he assigns TRUE to those propositions that provide an ordering rule within the domain of the problem, that is, to the most general predicates. Although such decisions alter the content of S's knowledge, they do not necessarily alter its form (i.e., its hierarchical organization). This formulation corresponds to the philosophers' model of counterfactual reasoning, where the reasoner is said to (a) construct a possible world in which the assumption can be true and (b) order available descriptions about that world in terms of their ability to structure to domain, that is, in terms of modal categories (Rescher, 1961, 1964).

The generality coding model would predict that, for class-inclusion relations, where information is readily organized hierarchically, Ss would assign generalities TRUE more often than chance. This is because generalities provide a convenient rule of inference; they define the domain much as laws and definitions in science do. Reasoning from generalities permits the imposition of a hierarchical organization on a set of elements that was disorganized by a counterfactual assumption.

For property-assignment relations, the generality coding model would also predict a marked preference for assigning generalities TRUE to the extent that property relations can be represented hierarchically. Recent work on semantic memory by Collins and Quillian (1969, 1970a, b) argues that properties, like the classes to which they belong, are hierarchically organized. That is, properties are sometimes stored in memory at the most superordinate node possible. For example, the properties of *snakes* (e.g., cold bloodedness) that are also shared by the superordinate, *reptiles* would be stored in long-term memory with *reptiles* and not with *snakes*. Therefore, property

relations should be treated in the same manner as class relations. The present data show that they are not. Although there is no significant difference in the overall performance of reasoners on class-inclusion and property-assignment problems, the pattern of responses of the two groups differs. Therefore, if generalities are assigned TRUE because of their organizing function, S's representation of property-assignment information must be different from class-inclusion information. This is precisely what was shown in a test of Collins and Quillian's hypothesis by Smith, Haviland, Buckley, and Sack (in press). In that study, Ss judged the truth or falsity of property-assignment relations in an artificially constructed hierarchy. The latency data indicated that all properties of a noun class are stored with—and retrieved directly from—that noun class. Consequently, in the present experiments, when Ss fail to reason with generalities, they do so on just those problems where there is independent evidence that information is not normally represented hierarchically. Considering the findings of Smith et al (in press), it is noteworthy that in the present experiment a majority of the Ss used a generalist strategy when solving the property-assignment problems. This may attest to the power of generalities for facilitating the ordered representation of information.

The generality coding model of S's knowledge predicts not only his acceptance of generalities but also his differential preference for assigning affirmative and negative generalities TRUE. In terms of a relative ability to provide an ordering rule, negative generalities only tell in which classes an object is not contained and leave uncertain the class in which it is contained. Consequently, the difference in the truth-value assignments to affirmative and negative generalities may be a result of the greater structuring of information provided by affirmative generalities. The present model would predict that (a) when the affirmative and negative statements provide equivalent structural information, their truth-value assignments will be equivalent and (b) the difference in preference for affirmatives and negatives should increase as the uncertainty of the domain increases (i.e., from a binary domain to an indefinitely distributed one).

The present data support the contention that Ss prefer to reason with generalities rather than with particulars. There is also considerable evidence that Ss will force information into an ordering described by general predicates, even when no such predicates have been provided (in contrast to the present paradigm). For example, Dawes (1964, 1966) showed that, when remembering a story, Ss erroneously recall particular statements (e.g., *Some As are Bs*) as universal ones (e.g., *All As are Bs*), a process Dawes calls "overgeneralization." His data show that such overgeneralizations were more frequent than erroneously recalling universal statements as particular ones, a process called "pseudodiscrimination." Dawes's findings

are of interest here in two ways. First, both the memory and prevaricative reasoning tasks can be viewed as inference tasks (cf. Fredericksen, 1972). Second, preference for generalities in reasoning may be related to a tendency to miscode particular propositions as generalities in memory. As a preliminary test of such a relationship, the Dawes memory task was administered to the 24 Ss in the reasoning condition of Experiment II. The data replicated the original findings: When Ss misrecalled the information in the paragraph, they overgeneralized more often than they pseudodiscriminated ($p < .001$). Of special interest is that the frequency of assigning affirmative generalities TRUE in the reasoning task is related to the tendency to overgeneralize in the memory task ($r = .68$, $p < .001$). This relationship between reasoning and miscoding was pointed out by Henle (1962), who argued that many errors in formal syllogistic reasoning tasks could be attributed to S's misinterpretation of the propositions—in a manner similar to the way the present Ss misinterpreted the information in the paragraph.

Therefore, the processes shown to operate in prevaricative reasoning may have commonalities with those shown for syllogistic reasoning and memory for paragraphs. What is needed is the further development of a model, perhaps along the lines of the generality coding model, that is sufficiently broad to account for inferences on both reasoning and memorial tasks.

REFERENCES

- Bartlett, F. *Remembering: A study in experimental and social psychology*. Cambridge: Cambridge University Press, 1932.
- Collins, A., & Quillian, M. Retrieval time from semantic memory. *Journal of Verbal Learning & Verbal Behavior*, 1969, 8, 240-248.
- Collins, A., & Quillian, M. Experiments on semantic memory and language comprehension. In L. Gregg (Ed.), *Cognition in learning and memory*. New York: Wiley, 1970a.
- Collins, A., & Quillian, M. Facilitating retrieval from semantic memory: The effect of repeating part of an inference. *Acta Psychologica*, 1970b, 33, 304-314.
- Dawes, R. Cognitive distortion. *Psychological Reports*, 1964, 14, 443-459.
- Dawes, R. Memory and distortion of meaningful written material. *British Journal of Psychology*, 1966, 57, 77-86.
- Fredericksen, C. H. Effects of task induced cognitive operations on comprehension and memory processes. In R. Freedle and J. B. Carroll (Eds.), *Language comprehension and the acquisition of knowledge*. New York: Winston, 1972.
- Henle, M. On the relation between logic and thinking. *Psychological Review*, 1962, 69, 366-378.
- Kintsch, W. Notes on the structure of semantic memory. In E. Tulving and W. Donaldson (Eds.), *Organization and memory*. New York: Academic Press, 1972.
- Meyer, D. On the representation and retrieval of stored semantic information. *Cognitive Psychology*, 1970, 1, 242-299.
- Myers, J. *Fundamentals of experimental design*. Boston: Allyn & Bacon, 1966.
- Random House dictionary of the English language*. (1st ed.) New York: Random House, 1966.
- Rescher, N. Belief-contravening suppositions. *Philosophical Review*, 1961, 70, 176-195.

- Rescher, N. *Hypothetical reasoning*. Amsterdam: North Holland, 1964.
- Revlis, R., & Hayes, J. R. The primacy of generalities in hypothetical reasoning. *Cognitive Psychology*, 1972, 3, 268-290.
- Revlis, R., Lipkin, S., & Hayes, J. The importance of universal quantifiers in a hypothetical reasoning task. *Journal of Verbal Learning & Verbal Behavior*, 1971, 10, 86-91.
- Simon, H. On the definition of causal relation. *Journal of Philosophy*, 1952, 49, 517-528.
- Simon, H. The logic of rational decision. *British Journal for the Philosophy of Science*, 1965, 16, 169-181.
- Simon, H., & Rescher, N. Cause and counterfactual. *Philosophy of Science*, 1966, 33, 323-340.
- Smith, E., Haviland, S., Buckley, P., & Sack, M. Retrieval of artificial facts from long-term memory. *Journal of Verbal Learning & Verbal Behavior*, in press.
- Steinberg, D. Negation, analyticity, amphigory, and the semantic interpretation of sentences. *Journal of Experimental Psychology*, 1970, 84, 417-423.
- Wilkins, M. The effect of changed material on the ability to do formal syllogistic reasoning. *Archives of Psychology*, 1928, 102, 1-83.

NOTES

1. Van Hook, K. J., Rose, B. J., & Taplin, J. E. Processes in hypothetical reasoning. In preparation.
2. Although the correlation measures make the first test redundant, the latter is included here in the interest of clarity.

(Received for publication May 21, 1973;
revision received August 17, 1973.)