JOURNAL OF VERBAL LEARNING AND VERBAL BEHAVIOR 20, 540–551 (1981)

From Words to Meaning: A Semantic Illusion

THOMAS D. ERICKSON

University of California, San Diego

AND

MARK E. MATTSON

State University of New York at Stony Brook

How are the meanings of individual words combined to form a more global description of meaning? This paper describes a phenomenon which sheds some light on one aspect of this process. Consider the following question: How many animals of each kind did Moses take on the Ark? Most people answer "two" even though they know quite well that it was Noah and not Moses who sailed the Ark. This illusion occurs even when time pressure is eliminated and subjects are told that questions may be "wrong" and given an example of a question with an inconsistent name in it. Two explanations of this illusion—that people are skipping over the name or that the focus of the question is leading them astray—are eliminated. Results indicate that it is important that the inconsistent name share semantic features with the correct name. An explanation of the illusion is developed.

A central process in language comprehension is the construction of a global description of the sentence meaning from the meanings of the individual words which make up the sentence. To begin with, we will make two assumptions: First, this construction process presumably involves discovering what relations hold between the meanings of the individual words. Thus, in the sentence "George Washington threw the coin across the Delaware," comprehension involves knowing that George Washington was the agent who performed the action of throwing and so on. Our second assumption is that the concepts or meanings represented by words are often complex. For instance, the verb "throw" refers not only to the physical action, but also implies the existence of an agent performing the action and of an object being thrown. And familiar proper nouns—such as "George Washington"—refer to a whole body of facts and scenarios: first president, general, crossing the Delaware, and so on. In discussing the details of the mechanism by which word meanings are combined, the complexity of the concepts referred to by words requires that our discussion be at the level below that of word meaning. So, for expository purposes, we will speak of words—or the concepts represented by words—as being decomposable into sets of semantic features.

Given that words consist of sets of semantic features, what do we know about ways in which the semantic features of one word might be connected up to those of other words? That is, how might the connecting relations be determined. Although we know very little about how this *actually* occurs in language comprehension, there is a good deal of research, largely in the do-

This research was supported by NSF Grant BNS79-24062 to James L. McClelland and NSF Grant BNS79-06024 to Bernard J. Barrs. In addition, the research was carried out while the first author was supported by an NSF Graduate Fellowship, and the second author by a grant to the Program in Cognitive Science from the Alfred P. Sloan foundation. The authors would like to thank the members of the LNR research group, particularly Jay McClelland, and also Bernie Barrs and Jean Mandler for helpful suggestions at various stages of this project. Reprint requests may be sent to Thomas Erickson, Psychology (C-009), UCSD, La Jolla, Calif. 92093.

0022-5371/81/050540-12$02.00/0

main of artificial intelligence, on how it *could* occur. One way is to make use of the constraints within the meanings of the words themselves. This approach is exemplified in the conceptual dependency theory of Schank and his colleagues (Schank, 1972; Riesbeck & Schank, 1978). A second way is to use world knowledge to determine the relations between words (e.g., Bransford & Johnson, 1972; Schank & Abelson, 1977; Winograd, 1972). Perhaps because this research has shown that creating computer programs which build global descriptions of meaning from individual words is very difficult, it has become widely assumed that anything which can contribute to the efficacy of the comprehension process, actually does (e.g., Just & Carpenter, 1980). In other words, it has become widely assumed that sentences are subject to exhaustive analysis and consistency checks during processing. But this is not the case. People do not always understand what is said to them; sometimes they fail to understand, sometimes they misunderstand. And while these failures of comprehension are sometimes due to lack of appropriate knowledge or error on the part of the speaker, there are other cases in which such failures occur when the understander possesses all the knowledge necessary for correct understanding. This paper explores such a phenomenon.

In response to the question, "How many animals of each kind did Moses take on the Ark?" most people answer "two." This is rather unusual, for it occurs even when they know, with complete certainty, that it was Noah who built and sailed the Ark. Although the question is really meaningless because of the inconsistent name, people do not notice the inconsistency. The explanation is not that people are simply being compliant. For one thing, the people who answer this question express surprise when the inconsistent name is pointed out. And for another, if people are given a series of questions and told to answer "wrong" when the question contains inconsistent information, many of them still make the error. Clearly something odd is occurring here.

The goal of this research is to document what we have come to call the Moses illusion, to develop an explanation for it, and to discuss its implications. We will begin by considering a rather trivial explanation of the illusion: Victims of the illusion often assert that they skipped over the name—not even seeing it. They point out, quite correctly, that the question contains redundant information and can be answered without knowing what the name is (How many animals were taken on the Ark?). In other words, according to this explanation the inconsistent name is never encoded—no phonemic or graphemic trace is ever laid down—and thus, from the perspective of the information processor, there is no inconsistent name to be processed and detected.

Experiment 1

In this experiment people are asked to read each question out loud before answering it, thus ensuring that the inconsistent name is encoded at at least a phonemic level. If a failure to encode the inconsistent name lies at the root of the illusion then this manipulation should eliminate it.

A secondary goal of this experiment is to increase the generality of the Moses illusion. To this end an attempt was made to create other questions which would act in the same way as the Ark question. The new questions were generated using largely intuitive criteria suggested by the Ark question; the Ark question itself was encountered in a book of puzzles (Fixx, 1977). On the basis of some pilot work, three other questions were chosen which either showed the strongest tendency toward producing the illusion or were interesting for some other reason. Table 1 shows the four stimulus questions.

Method

Subjects. The subjects were 28 undergraduates drawn from lower-division psychology courses at the University of

TABLE 1
TARGET QUESTIONS

How many animals of each kind did Moses take on the Ark?
In the biblical story, what was Joshua swallowed by?
What is the nationality of Thomas Edison, inventor of the telephone?
In the novel "Moby Dick," what colour was the whale that Captain Nemo was after?

California at San Diego who participated in the experiment for course credit. All subjects were native speakers of English.

Apparatus. A Cromenco Z-2 microprocessor interfaced to an ACT-IV terminal with a DEC black and white monitor was used to present the stimuli. A Teac A-170S cassette deck and a Realistic condenser microphone were used to record the subjects' responses on the verbal task.

Procedure. The subject was seated in front of the terminal, on which the first set of instructions appeared.

> The first section of this experiment involves answering questions. The questions will appear one at a time on this screen. After a short time they will disappear. Your job is to read the question out loud, then answer it aloud as fast as you can. With your permission we will tape your answers. Usually a one word answer will be enough. If you don't know the answer or can't recall it right away, just say "don't know." You will occasionally encounter a question which has something wrong with it. For example, you might see the question: "Why was President Gerald Ford forced to resign his office?" The thing that is wrong in this example is that Ford wasn't forced to resign. When you see a question like this, just say "wrong" (or something similar).
>
> So, in summary: you will see a question. Read it out loud, then answer it as quickly as possible. If you don't know the answer, say "don't know." Or if the question is wrong, say "wrong." Following these instructions are three practice questions.

When the subject was ready to continue, the three practice questions were presented, each followed by its correct answer and justification. One of these questions contained an inconsistent name, with the answer being given as "wrong" and the justification citing the inconsistent nature of the name. After this, any questions the subject had were answered, and the first section of the experiment was begun.

Each of the 20 stimulus questions was presented for 5 seconds, centered on an otherwise blank screen. The presentation of each question was triggered by the subject. The questions were presented in the same order for each subject. This part generally lasted from 3 to 5 minutes.

Immediately after the end of the first part of the experiment, the second section was begun. Subjects were told they would be tested for their memory of some of the questions they had just seen. They were presented with the first part (up to, but not including the name) of each target question, and asked to write down the rest of the question "as you remember seeing it." The subjects were told that if they could not recall the answer fairly soon, to write "can't recall," and go on to the next item.

The third section of the experiment was a check of the subjects' knowledge. The subjects were presented with four questions which directly probed their knowledge of the agent in each of the target questions (e.g., "Who was it that took the animals on the Ark?").

Stimuli. In addition to the target questions, 16 distractors were constructed. Some distractors were simple questions, some were more difficult, and 3 were obviously wrong. These obviously wrong questions contained inconsistent names which were always detected by the subjects. For example: "When he stepped on the surface of the moon, what were Ronald Reagan's words?" These questions were included to demonstrate that the subjects knew what to do when faced with an inconsistent question.

Results

Two types of results are reported: the percentage of times the illusion occurred and the percentage of memory deviations. The illusion was judged to have occurred when the subject answered a target ques-

tion as though it made sense[1] (e.g., answering "two" to the Ark question, instead of saying "wrong" or "don't know") even though the subject had the correct knowledge (e.g., knew that Noah sailed the Ark). Thus the percentage of occurrences of the illusion is based on the number of subjects who had the correct knowledge (see the last column of Table 2), and not on the total number of subjects run. A memory deviation was an occurrence of the illusion followed by an incorrect recall of the target question as containing the name that would have been correct instead of the inconsistent name (e.g., recalling the Ark question as having had the name Noah in it). Occurrence of a memory deviation is especially strong evidence that the subject had the correct knowledge and that the knowledge was readily available. As before, the percentage of memory deviations is based on the number of people having the correct knowledge.

The results may be seen in Table 2. The Moses illusion occurred for all four questions; the frequency of the illusion differed significantly across the four questions, $\chi^2(3) = 8.91$, $p < .05$. Memory deviations occurred for three of the four questions.

TABLE 2
RESULTS FROM EXPERIMENT 1

Question	Frequency of illusion (%)	Memory deviations (%)	Correct knowledge
Ark	81	73	26
Whale	40	10	10
Telephone	44	0	16
Moby Dick	44	33	9

Note. Although there were 27 subjects in each group, the percentages were based on the number of people having the correct knowledge (column 3).

Discussion

The important finding is that the Moses illusion occurs even when the subject reads the question aloud, thus ensuring that the inconsistent name is encoded. This is in line with an observation we made during our pilot work: namely, that when subjects were asked to repeat the Ark question immediately after answering it, they usually repeated veridically: with Moses, not Noah. Secondly, note that the Moses illusion occurs for each of the four questions, thus indicating that the illusion has at least some generality. And finally note that the illusion occurred even though the subjects were told that some questions might be "wrong" and given an example of a question with an inconsistent name in it.

Since the trivial explanation can be rejected, let us consider a second possible explanation of the Moses illusion. The second explanation has to do with the fact that the Moses illusion involves questions. Perhaps the illusion occurs because the focus of each question is on something other than the inconsistent name (e.g., for the Ark question the focus is on "how many animals" and not on "who").

EXPERIMENT 2

Experiment 2 addresses the question of whether the Moses illusion is solely due to the effects of the focus of the question. In this experiment the target questions from Experiment 1 are transformed into statements (see Table 3) which the subjects are to judge as true or false. If the Moses illusion is due solely to a misdirection of focus, people should have no trouble in answering false to the statement: "Moses took two animals of each kind on the Ark." On the other hand, if the Moses illusion is a result of some more general process, then some subjects should answer true to the Moses

[1] This does not mean the answers were necessarily correct. For instance, a number of subjects answered the Moby Dick question by saying that the whale was gray. On the other hand, answers of "don't know" were not counted as occurrences of the illusion because of the possibility that—in spite of instructions and examples—subjects were unsure about when to say a question was wrong. And finally, cases where the subjects answered the question and then immediately changes their answer were not counted as occurrences of the illusion.

TABLE 3
TARGET STIMULI FOR EXPERIMENT 2

Moses took two animals of each kind on the Ark.
In the biblical story, Joshua was swallowed by a whale.
Thomas Edison, inventor of the telephone, was an American.
In the novel Moby Dick, Captain Nemo was chasing a white whale.

statement in spite of knowing that Noah built and sailed the Ark. Thus the result of interest is simply whether or not the illusion will occur with statements.

Method

Subjects. The subjects were 25 undergraduates from lower-division psychology courses at the University of California at San Diego, who participated in the experiment for course credit. The subjects were all native speakers of English.

Stimuli. In addition to the four target statements, the 16 distractor questions from Experiment 1 were transformed into true or false statements.

Materials. The materials for this experiment consisted of test booklets containing directions and stimuli. Part 1 of the booklet consisted of directions, followed by three example statements with answers and justifications, followed by 20 stimuli, one on each page. The order was randomized for each booklet, with the constraint that two target statements should not appear consecutively. Parts 2 and 3 of the booklet consisted of a memory test and knowledge test analogous to those in Experiment 1.

Procedure. Subjects were run in groups ranging in size from 2 to 6. Subjects were each given a test booklet and allowed to proceed at their own pace (although the directions did emphasize that they should try and respond quickly). In brief, the directions on the booklet told the subjects that their job was to read a statement and circle either "true," "false," or "don't know"; as an example they were told that the statement "Former president Gerald Ford was forced to resign his office" was false. The directions also emphasized that subjects should not go back to previous statements.

Generally subjects were able to go through the booklet within about 5 minutes.

Results

The results may be seen in Table 4. The Moses illusion occurred for each of the four target statements. There was no detectable difference in the frequency of occurrence of the illusion across the different statements, χ^2 (3) = 2.58, $p > .05$. Memory deivations occurred for three of the four target statements; the exception being the Moby Dick statement.

Discussion

There are three things to note here: First, the Moses illusion is not solely due to effects of the focus of the questions since it continues to occur for each of the four statements. Second, note that in transforming the Ark and Telephone questions into statements, in addition to the change in focus, there was also a change in the relationship of the inconsistent name to the rest of the sentence: namely, for these two statements, the inconsistent name was the first thing in the sentence. Since the illusion occurred for these statements as well as the others, it seems that the inconsistent name does not have to be preceded by a biasing context. Finally, unlike the previous ex-

TABLE 4
RESULTS FROM EXPERIMENT 2

Statement	Frequency of illusion (%)	Memory deviations (%)	Correct knowledge
Ark	41	41	22
Whale	27	9	11
Telephone	26	5	19
Moby Dick	12	0	8

Note. Although there were 25 subjects in each group, the percentages are based on the number of subjects having the correct knowledge (column 3).

periment, there were no actual limits on reading time: subjects read through the booklet at their own pace, and yet still made the error. Experiment 2 thus increases our confidence in the robustness of the illusion by showing that the illusion occurs in the absence of time pressure and that the illusion occurs for statements as well as for questions.

Experiment 3

Experiment 3 examines the role of the inconsistent name in the Moses illusion. In thinking about the Moses illusion, it seems evident that there are some names which subjects would always notice. For example, it seems unlikely that subjects would fail to notice that something was wrong with the question: "How many animals of each kind did George Washington take on the Ark?" Accepting this assumption for the moment, the question then becomes: In what ways must the inconsistent name be related to the correct name in order for it to produce the illusion? If we consider the Ark question and the relations between Moses and Noah, two levels of similarity become apparent. For one thing, the names Moses and Noah are phonologically similar: both names have two syllables with the first stressed; and the initial syllables are also quite similar in sound, each being a nasal followed by a long "o." At a higher level Moses and Noah can be viewed as being similar because they share a number of semantic features—they are main figures in biblical stories, they received messages from God, lead their followers out of adversity, and so on. In Experiment 3 we look at the evidence for each of these hypotheses by manipulating the characteristics of the inconsistent names which appear in the target questions.

Two hypotheses are examined in this experiment. (1) Phonological similarity: The more phonologically similar an inconsistent name is to the correct name (other things being equal), the more likely it is that the Moses illusion will occur. (2) Semantic similarity: The more semantically similar an inconsistent name is to the correct name, the more likely it is that the illusion will occur. To test the phonological similarity hypothesis, the names Moses/Joshua, Adam, and Abraham were inserted in the Ark/Whale question (see Table 5). There were two criteria for selecting these inconsistent names. One was that the names be as similar at the semantic level as possible; this was achieved by constraining the names to the set of Old Testament biblical figures. The second criterion was that the names have varying degrees of phonological similarity, where phonological similarity was judged by the following features: same initial sound, same terminal sound, same syllable stressed, and same number of syllables.

The semantic similarity hypothesis was tested in all four questions. In the Ark and Whale questions it was tested by inserting a nonbiblical inconsistent name in each question, with the prediction of the semantic similarity hypothesis being that the illusion would not occur for either nonbiblical name. The name Nixon was used in the Ark question, and the name Jeffrey in the whale question (see Table 5). The name Nixon was used because, like the other names in the Ark question, it was judged to have a fairly rich set of semantic units associated with it. The name Jeffrey was used because it was judged to have few semantic units associated with it. While the failure of both names to produce the illusion would provide support for the semantic similarity hypothesis, if only the name Nixon failed to produce the illusion, then this would suggest something somewhat different: namely, that the illusion required only the absence of semantic dissimilarity, and not the presence of semantic similarity to occur. In other words, the issue is, if the name Nixon fails to produce the illusion, does it fail because it lacks semantic features which are similar to those of Noah or does it fail because it has semantic units which are very different from those of Noah?

The semantic similarity hypothesis was also tested using the Telephone and Moby

TABLE 5
INCONSISTENT NAMES USED TO TEST THE VARIOUS HYPOTHESES IN EXPERIMENT 3

Hypothesis tested	Questions used	Names used in questions
Phonological similarity	Ark	Moses
		Adam
		Abraham
	Whale	Joshua
		Adam
		Abraham
Semantic similarity	Ark	Nixon (versus biblical names)
	Whale	Jeffrey (versus biblical names)
	Telephone	Thomas Edison
		Benjamin Franklin
		George Washington
	Moby Dick	Nemo
		Jonah
		Nelson

Dick questions. Two tacks were taken: In the case of the Telephone question, three names—Thomas Edison, Benjamin Franklin, and George Washington—were selected because they seemed to vary in their degree of semantic similarity to the target. In the case of the Moby Dick question, the names Nemo, Jonah, and Nelson were selected;[2] although each had semantic similarities to the target, the semantic features were different ones. For example, although Jonah and Ahab were both involved with whales, and Nemo and Ahab were both captains of seagoing craft, it is much less evident that Jonah and Nemo share semantic features.

Method

Subjects. The subjects were 187 members of an introductory psychology class who participated in the experiment as part of an in-class demonstration.

[2] The third name, Nelson, was originally intended to play a role similar to that of Jeffrey, as a name with few associations. The authors later realized that Nelson was the name of the 19th-century British admiral; and it was also pointed out that there was a Captain Nelson on a TV series entitled "Voyage to the Bottom of the Sea." So, inadvertently, the third name also had nautical associations.

Stimuli. The four target questions were the same as in Experiment 1 (see Table 1), except that, in this experiment, different versions of each question were formed by inserting one of three or four different inconsistent names in each question (see Table 5).

Materials. Test booklets were made up. The booklets were identical to those used in Experiment 2 except for the following changes: the stimuli were questions (from experiment 1) instead of statements; there were appropriate changes in the directions; and, since some versions of different questions used the same inconsistent name (see Table 5), booklets were constructed so that no inconsistent name was used more than once per booklet.

Procedure. The procedure was identical to that used in Experiment 2, except that the experiment was run, in class, on all the subjects at once.

Results

The results of Experiment 3 appear in Table 6. The phonological similarity hypothesis was tested by inserting three inconsistent names in the Ark and Whale questions. Evidence in favor of this hypothesis would be increasing frequencies of

the illusion, with the lowest for Abraham, the next higher for Adam, and the highest for Moses/Joshua. As can be seen in Table 6 this pattern did not occur for either question. Instead there were either no significant differences among the three names, as in the case of the Ark question, $\chi^2(2) = 5.75$ $p > .05$, or, as in the case of the Whale question, there were significant differences, $\chi^2(2) = 12.20$, $p < .01$, but the pattern of percentages was not as predicted.

In the Ark and Whale questions the semantic similarity hypothesis was tested by comparing the frequency of the illusion for the biblical names with the frequency of the illusion for the nonbiblical name. For both the Ark and Whale questions there were significant differences in frequency of the illusion for biblical versus nonbiblical names, $\chi^2(1) = 36.02$, $p < .01$, and $\chi^2(1) = 7.02$, $p < .01$, respectively. Thus, there do seem to be differences between biblical and nonbiblical names for the two questions. On the other hand, note that there are also the differences within the set of biblical names for the Whale question (Table 6); this will be considered in the discussion.

The semantic similarity hypothesis was also tested using the Telephone and Moby Dick questions. In the Telephone question the inconsistent names were chosen so as to differ in their semantic similarity to the correct name, the prediction being that the more semantically similar the name, the more often the illusion should occur. This pattern of results was obtained. Comparisons between the frequencies of the illusion for the names Edison and Franklin and between the frequencies of the illusion for the names Franklin and Washington showed significant differences for both cases, $\chi^2(1) = 5.75$, $p < .05$, and $p < .01$ according to the Fisher exact probability test,[3] respectively. Thus, the results for the Telephone question also provide support for the semantic similarity hypothesis.

[3] Whenever the Fisher exact probability test is used, it is because the χ^2 test could not be used because the expected frequencies were too low.

And finally, the results from the Moby Dick question also supported the semantic similarity hypothesis. As Table 6 indicates, the illusion occurred with fairly high frequency for each of the three names; additionally, there was no detectable difference in frequency of the illusion for any of the three names, $\chi^2(2) = .481$, $p > .05$. Unlike the sets of names which produced the illusion in the previous questions—Moses/Adam/Abraham and Edison/Franklin—which are similar both to the correct name and to each other (all, for instance, could be classed as members of the categories of biblical figures or inventors), the names Nemo, Jonah, and Nelson seem to have few, if any, similarities to each other, thus suggesting that no special type of similarity (such as having the same superset) is necessary for the illusion.

Discussion

The phonological similarity hypothesis did not hold up; the predicted pattern did not appear for either the Ark or Whale questions. However, given the small number of items, the roughness of the phonological similarity metric, and the difficulty of varying phonological features while holding semantic similarity constant, the possibility of an effect due to phonological similarity cannot be entirely discounted. It is also possible that in spite of the attempt to hold semantic similarity constant, it varied enough to obscure any effects of phonological similarity.

The semantic similarity hypothesis fared better: it received support in each of the four questions. For both the Ark and Whale questions there were significant differences between the frequency of the illusion for biblical names and the frequency of the illusion for nonbiblical names. However, while it is tempting to conclude that these differences are due to the category manipulation, it should also be noted that an overall difference did occur for the biblical names used in the Whale question. In this question there was a substantial effect for Joshua, a

TABLE 6
RESULTS FROM EXPERIMENT 3

Question	Frequency of illusion (%)	Memory deviations (%)	Correct knowledge	Total number
Ark (Moses)	49	44	39	42
Ark (Adam)	70	61	33	36
Ark (Abraham)	44	43	63	67
Ark (Nixon)	0	0	42	42
Whale (Joshua)	39	15	33	68
Whale (Adam)	0	0	21	41
Whale (Abraham)	17	8	24	36
Whale (Jeffrey)	0	0	27	42
Telephone (Edison)	46	0	35	51
Telephone (Franklin)	20	0	44	57
Telephone (Washington)	1	0	72	80
Moby (Nemo)	58	31	26	79
Moby (Jonah)	48	12	25	57
Moby (Nelson)	53	13	15	51

Note. The percentages were based on the number of subjects having the correct knowledge (column 3).

smaller one for Abraham, but none for Adam. The lack of an effect for Adam is surprising; this is not due to some special characteristic of the name Adam (such as high familiarity), since Adam induced the illusion most frequently (70% of the time) in the Ark question. Nor does it seem likely that this is due to some effect of the context of the inconsistent name (such as the question failing to specify the scenario referred to) since the illusion occurred when the question had the names Joshua and Abraham in it. A more likely possibility is that the absence of the illusion is due to an interaction between the name Adam and its context: perhaps the verb "swallowed" in conjunction with the "Adam" made subjects think of the Adam-eating-the-apple scenario, and thus lead them to notice the question was wrong. There is no apparent explanation of the absence of the illusion for the name Adam in the Whale question which would discredit the interpretation of biblical versus nonbiblical names difference as support for the semantic similarity hypothesis.

The semantic similarity hypothesis also received support from the Telephone and Moby Dick questions. In the Telephone question, the more semantically similar the inconsistent name was to the correct name, the more frequently the illusion occurred. And in the Moby Dick question, the names Nemo, Jonah, and Nelson each induced the illusion with regularity, although there seems to be no semantic feature common to all of them. This suggests that there is no special type of semantic feature which must be shared with the correct name for the illusion to occur. Thus, with the exception of the name Adam for the Whale question, each prediction made by the semantic similarity hypothesis was confirmed.

Finally, a few words about the memory deviations. In general, the frequency of memory deviations tended to parallel the frequency of the illusion, albeit at a somewhat lower level. The one exception to this is the Telephone question (What is the nationality of X (e.g., Thomas Edison), inventor of the telephone?), in which there was only one memory deviation out of 38 occurrences of the illusion. One possiblity is that, for some reason, the subjects had

always become aware of inconsistent nature of the name by the time they were asked to recall the Telephone question. In pilot work this was occasionally observed to occur with the Ark question. A second possibility is, assuming the occurrence of memory deviations indicates that the knowledge in question is highly accessible, perhaps the subjects did not have readily accessible knowledge about Bell inventing the telephone, and only recovered it when explicitly questioned about it. A third factor, although rather small, was that a few subjects recalled the question as being "of Thomas Edison, inventor of the light bulb."

General Discussion

Let us begin by summarizing what has been learned about the Moses illusion. The Moses illusion will occur when subjects read a target question—either silently or out loud—or hear the target question (latter observed during pilot work). Thus the illusion seems to be independent of the way the question is presented. Indeed, as pilot work showed, the illusion is sufficiently robust that subjects are able to repeat the question veridically without noticing anything wrong. Nor is the Moses illusion dependent on misdirection of focus due to the stimulus being a question; it will also occur in a statement. As Experiments 2 and 3 showed, the Moses illusion will occur in the absence of time pressure; and it occurs even when people are told to watch for questions which are wrong, and given examples of such questions (Experiments 1 and 3). And finally, as the Ark and Telephone statements in Experiment 2 showed, the inconsistent name need not be preceded by biasing context—the illusion will occur even when the inconsistent name is the first thing encountered by the subject. In short, the Moses illusion is resistant to a wide variety of manipulations of the context of the inconsistent name.

The manipulations which did make a difference in the frequency of the Moses illusion were manipulations involving the nature of the inconsistent name (Experiment 3). There are some inconsistent names for which the illusion never occurs (e.g., Nixon in the Ark question), and other inconsistent names which induce the illusion readily. For the Moses illusion to occur, it is necessary that the inconsistent name share some semantic features with the correct name, although there does not seem to be any particular semantic feature which must be shared (Moby Dick question). The results for the Telephone question in Experiment 3, while not definitive, suggest that the more semantic features shared, the more frequently the illusion will occur. The other aspect of the inconsistent name which was looked at, its phonological similarity to the correct name, did not appear to facilitate the occurrence of the illusion.

Now let us consider some explanations of what is going on. First, to provide a framework for this discussion, let us divide sentence comprehension into three parts: an encoding process, in which the stimulus is represented—but not identified—in the processing system; a lexical access process, which produces a set of semantic features which specify the meaning of each word; and a construction process, in which the semantic features of the individual words are combined with the semantic features of the other words to produce an overall description of the meaning of the sentence. Experiment 1 eliminated the possibility that the Moses illusion is due to a failure in the encoding process: the inconsistent name does get encoded.

Another possibility is that the illusion is a result of the lexical access process: that is, the lexical access process, rather than producing the semantic features of Moses, instead produces the semantic features of Noah. This would require arguing that the lexical access process was being influenced by the context of the name. However, inspection of Tables 1 and 3 reveals that even by the laxest criteria only three of the eight stimuli uniquely constrain the inconsistent name (the Ark question, and the Moby Dick

question and statement). Indeed, in Experiment 2, two of the stimuli have the inconsistent name with no preceding context (e.g., Moses took two animals of each kind on the Ark). Thus, not only would one be in the position of arguing that lexical access is influenced by context, a moot point (Mehler, Segui, & Carey, 1978; Swinney, 1979), but one would have to argue that lexical access of the inconsistent name was suspended until after lexical access of subsequent words—a position which would seem to be untenable.

If the illusion is not a result of the lexical access process, then it must be due to the way in which the semantic features of the individual words are put together to produce a description of the meaning of the sentence. Let us assume that the construction process begins with several sets of semantic features, each set representing the meaning of a word. The main thing the construction process must do in producing a description of the sentence is to establish what kind of relations hold among the semantic features of different words in the sentence. The question of interest is how this can occur when certain semantic features are inconsistent with the overall meaning of the sentence. Let us examine the way in which the construction process might work for a normal, consistent sentence.

Suppose we have the sentence "Moses took the Ten Commandments down from the mountain." Assuming that the construction process begins with sets of semantic features produced by the lexical access process, there will be a set of semantic features for Moses, such as human, male, biblical character, got Ten Commandments, confronted Pharaoh, led his people out of Egypt, another set of semantic features for took, and so on. The construction process will take these semantic features and determine what relations (many of which will be semantic features of the sentence itself) hold among them. Thus the knowledge that "Moses" refers to a human being and that human beings are capable of performing actions such as "taking things" (or conversely, that "take" generally requires a human agent), allows the determination of the relation between "Moses" and "took." Similarly, the knowledge that take generally requires a physical object and that the Ten Commandments were in the form of stone tablets, allows the determination of the relation between took and the Ten Commandments. However, the knowledge that Moses confronted Pharaoh is of no use in this sentence; there is no place in the description where it would make contact with a semantic feature of another word. But in this case, this failure of part of the meaning of "Moses" to make contact with the rest of the sentence is inconsequential; it does not indicate anything wrong with the sentence, it just indicates that one of the semantic units of "Moses" is not relevant to the description of the sentence. Since most nouns have large numbers of semantic features, some of which would fail to make contact in almost any given sentence, such an occurrence would not be unusual. Thus interrupting processing every time a semantic feature failed to fit would be disadvantageous. When would processing be interrupted? A more likely mechanism would be to interrupt processing only if there were very few or no semantic features of a word which fit, making construction of a complete description of the meaning of the sentence difficult or impossible.

Something like this seems to be at the root of the Moses illusion. As bundles of semantic features are connected to each other, semantic units which do not make contact with other units arouse no special note; it is only when a bundle is only sparsely (or not at all) connected to the description of the sentence that the construction process is terminated and the mistake is noticed. This conception of the construction process fits in nicely with the semantic similarity hypothesis since it says that the fewer semantic features of a word that fit

into the description (i.e., in these experiments, the more dissimilar the inconsistent word is to the correct word), the likelier it is that the inconsistency will be noticed.

Much of the above explanation is speculative. While we hope that future work will garner more evidence for these claims, the Moses illusion does illustrate two characteristics of language comprehension which have been generally ignored in theories of language comprehension: First, sentences are not subject to exhaustive analysis. Obvious errors—obvious both because subjects have the knowledge to recognize them as such and because they are told to watch for them—escape notice. And second, the comprehension process seems to be, in some sense, sensitive to partial meanings of words, since some inconsistent names will escape notice in a particular sentence, and others will not.

REFERENCES

BRANSFORD, J. D., & JOHNSON, M. K. Contextual prerequisites for understanding: Some investigations of comprehension and recall. *Journal of Verbal Learning and Verbal Behavior*, 1972, 11, 717–726.

FIXX, J. J. *More games for the superintelligent*. Popular Library: New York, 1977.

JUST, M. A., & CARPENTER, P. A. A theory of reading: From eye fixations to comprehension. *Psychological Review*, 1980, 87, 329–354.

MEHLER, J., SEGUI, J., & CAREY, P. Tails of words: Monitoring ambiguity. *Journal of Verbal Learning and Verbal Behavior*, 1978, 17, 29–35.

RIESBECK, C. K., & SCHANK, R. C. Comprehension by computer: Expectation-based analysis of sentences in context. In W. J. M. Levelt & G. B. Flores d'Arcais (Eds.), *Studies in the perception of language*. New York: Wiley, 1978.

SCHANK, R. C. Conceptual dependency: A theory of natural language understanding. *Cognitive Psychology*, 1972, 3, 552–631.

SCHANK, R. C., & ABELSON, R. P. *Scripts, plans, goals, and understanding*. Hillsdale, N. J.: Erlbaum, 1977.

SWINNEY, D. A. Lexical access during sentence comprehension: (Re)Consideration of context effects. *Journal of Verbal Learning and Verbal Behavior*, 1979, 18, 645–659.

WINOGRAD, T. *Understanding natural language*. New York: Academic Press, 1972.

(Received March 13, 1981)