

How Do We Know That We Know? The Accessibility Model of the Feeling of Knowing

Asher Koriat

How Do We Know That We Know? The Accessibility Model of the Feeling of Knowing

Asher Koriat

Even when Ss fail to recall a solicited target, they can provide feeling-of-knowing (FOK) judgments about its availability in memory. Most previous studies addressed the question of FOK accuracy, only a few examined how FOK itself is determined, and none asked how the processes assumed to underlie FOK also account for its accuracy. The present work examined all 3 questions within a unified model, with the aim of demystifying the FOK phenomenon. The model postulates that the computation of FOK is parasitic on the processes involved in attempting to retrieve the target, relying on the accessibility of pertinent information. It specifies the links between memory strength, accessibility of correct and incorrect information about the target, FOK judgments, and recognition memory. Evidence from 3 experiments is presented. The results challenge the view that FOK is based on a direct, privileged access to an internal monitor.

In their review of memory research, more than 20 years ago, Tulving and Madigan (1970) noted that "one of the truly unique characteristics of human memory [is] its knowledge of its own knowledge" (p. 477). They argued that

no extant conceptualization, be it based on stimulus-response associations or an information processing paradigm, makes provisions for the fact that the human memory system cannot only produce a learned response to an appropriate stimulus or retrieve a stored image, but it can also rather accurately estimate the likelihood of its success in doing it, (p. 477)

Although the last 20 years have engendered a large volume of research on the feeling of knowing (FOK), only a small proportion of this work has direct bearing on the question raised by Tulving and Madigan. In fact, the literature on memory monitoring displays an ambivalent attitude concerning people's ability to monitor their knowledge. Whereas some regard this ability as a mystery, others, mostly in the area of decision making, appear to take it for granted, focusing on explaining why sub-

jects' monitoring performance deviates from perfect accuracy (e.g., Lichtenstein, Fischhoff, & Phillips, 1982; See Koriat, Lichtenstein, & Fischhoff, 1980). Both approaches share a common attitude that discourages investigation of the processes underlying FOK accuracy. The present project instead aims precisely at explaining the basis for the accuracy of FOK, judgments in predicting future memory performance.

The FOK has attracted most attention in connection with memory blockage states, where subjects fail to retrieve the sought-after target but can nevertheless judge whether it is stored in memory or not. Such states, like the tip-of-the-tongue (TOT) state (see A. S. Brown, 1991; R. Brown & McNeil, 1966; Koriat & Lieblich, 1974), are puzzling because they combine two seemingly inconsistent features, the subjective conviction that one knows the solicited word or name and the actual inability to produce it. This discrepancy raises the question of how does a person know that he or she knows in the face of being unable to produce the solicited target? This question is the focus of the present study.

Three Questions Regarding the FOK

Three different questions regarding the FOK should be distinguished. Question 1 concerns the accuracy of FOK judgments, that is, their validity in predicting subsequent memory performance. A large number of studies confirmed that subjective conviction is predictive of objective memory performance. Thus, subjects unable to retrieve a solicited item from memory can estimate with above-chance success whether they will be able to recall it in the future, produce it in response to clues, or identify it among distractors (e.g., Freedman & Landauer, 1966; Gardiner, Craik, & Bleasdale, 1973; Gruneberg & Monks, 1974; Gruneberg & Sykes, 1978; Hart, 1965, 1967a, 1967b; Leonesio & Nelson, 1990; Nelson & Karens, 1990; Schacter, 1983). Other studies indicated that FOK judgments may also be effective as predictors of performance on implicit memory tasks (Lupker, Harbluk, & Patrick, 1991; Nelson, Gerler, & Narens,

A brief version of Experiment 1 was presented at the First International Conference on Memory, Lancaster, England, July 1991. This research was supported by the Israel Science Foundation, which is administered by the Israel Academy of Sciences and Humanities. The experiments were conducted at the Institute of Information Processing and Decision Making, University of Haifa, Haifa, Israel. The writing of this article was completed while Asher Koriat was the Irv Acenberg visiting scientist at the Rotman Research Institute of Baycrest Centre, Toronto, Ontario, Canada.

I wish to acknowledge my appreciation and debt to the late Israel Lieblich, who was a partner in my first steps in this area.

I want to thank Michai Sion for her help in all stages of the research program; Morris Goldsmith, Larry Jacoby, Janet Metcalfe, and Tom Nelson for their comments on an earlier draft; and Hagit Tal and Dahlia Mevorach for the programming of the experiments.

Correspondence concerning this article should be addressed to Asher Koriat, Department of Psychology, University of Haifa, Haifa, Israel.

1984; Yaniv & Meyer, 1987). The standard finding is that the predictive validity of FOK judgments is above chance, though far from perfect.

Question 2 pertains to the underlying process. The few studies that addressed this question have focused on delineating the basis of FOK judgments (e.g., Begg, Duft, Lalonde, Melnick, & Sanvito, 1989; Koriat & Lieblich, 1977; Leonesio & Nelson, 1990; Lupker et al., 1991; Nelson et al., 1984; Reder, 1987, 1988; Reder & Ritter, 1992; Schacter & Worling, 1985; Schwartz & Metcalfe, 1992). A comprehensive review of possible mechanisms underlying FOK judgments was provided by Nelson et al. (1984; see also Krinsky & Nelson, 1985).

Question 3 involves relating accuracy to process: An adequate process model of FOK must not only specify the determinants of FOK judgments but it must also indicate how the processes leading to FOK can account for both the success and failure of FOK in monitoring subsequent memory performance.

Most FOK studies have confined themselves to Question 1, only a few studies touched on Question 2, and practically none of these directly addressed Question 3. Furthermore, as I show later, many of the studies concerned with the underlying process fail to distinguish between the perspective of the experimenter and that of the subject, that is, between objective and subjective properties of memory. Also, some of the discussions of the FOK failed to distinguish between Questions 2 and 3 stated earlier, so that the question "how does one know that one knows?" is sometimes interpreted to mean "why are people successful in monitoring their knowledge?" (Question 3), and sometimes it simply means "what are the determinants of the FOK?" The reason for this confusion may lie in the term "knowing" (in contrast with the term "believing") being understood as implying correct knowledge (cf. the frequent use of the term "false belief" with the infrequent use of the term "false knowledge").

The present article represents an effort to deal with all three questions within a common framework, using Question 2 as the starting point. In brief, the model I detail here assumes that the computation of FOK is parasitic on the processes involved in attempting to retrieve the memory target, relying mostly on the amount of relevant information that is recruited. It is hoped that the model, in combination with the experimental findings to be reported, will contribute toward the demystification of the FOK phenomenon.

Two Approaches to the FOK

How does one know that one knows something that he or she is unable to retrieve? Two general approaches to this question may be distinguished. The first, trace-based view (see Nelson et al., 1984) assumes the existence of an *internal monitor* that directly detects the presence of the target in store, and it is this monitor that is consulted in making FOK judgments. This position is represented by Hart's approach in his pioneering studies of the FOK. (Hart, 1965, 1967a). Hart stressed the functional value of having such a monitor, given the general fallibility of the memory system (see also Yaniv & Meyer, 1987). In such a system, FOK

can serve as an indicator of what is stored in memory when the retrieval of a memory item is temporarily unsuccessful or in-

terrupted. If the indicator signals that an item is not in storage, then the system will not continue to expend useless effort and time at retrieval; instead, input can be sought that will put the item into storage. Or if the indicator signals that an item is in storage, then the system will avoid redundantly inputting information that is already possessed, (Hart, 1965, p. 214)

The assumption, then, is that subjects have direct access to the information pertaining to the presence of the solicited item in memory and that this information appears in a ready-made format.

At first sight, this solution to the question of how one knows that one knows appears to raise the homunculus problem of how the monitor itself can know. However, the idea becomes much less far fetched when, an analogy is drawn with the organization of information in a computer. The names of the files stored on a computer disk are normally listed in a separate directory file. Therefore, if I want to check whether a specific file exists on the disk, I need only consult the directory to see whether it contains the name of the file, without having to access the content of the file itself. In a similar manner, the internal monitor underlying FOK judgments may be thought of as a mechanism that has access to a directory that catalogs the files stored in memory, that is, to the listing of the "names" of the "files" stored. Clearly, such a directory is of great value because it can save the time and effort searching for something that is not there, just as it does in computerized information systems.

The internal-monitor view implies a two-stage conception of retrieval: When presented with a memory query, subjects first consult the internal monitor to check whether the sought-after target is stored in memory (analogous to consulting the directory listing in a computer) before embarking on an attempt to retrieve it (analogous to accessing the file itself; see, e.g., Figure S in Nelson & Narens, 1990; Reder, 1988). This two-stage conception assumes that FOK judgments rest on a process that is independent of the process required to retrieve the target itself. Of course, when subjects are asked to provide only FOK judgments (e.g., Reder & Ritter, 1992), they need only use the first memory-monitoring process.

The internal-monitor view is responsible, perhaps, for the paucity of experimental work on the determinants of FOK, because FOK judgments are implicitly conceived as the mere output of a specialized, encapsulated mental module (see Fodor, 1983). Furthermore, this view implies a conception of FOK as an all-or-none indicator, rather than as a graded signal. Indeed, in studies that focus on retrieval failures, it is typical to solicit FOK judgments in a discrete form, distinguishing between "positive" and "negative" FOK judgments, as if each refers to a different, discrete "memory state" (see Koriat & Goldsmith, 1993b).

The second approach posits that FOK judgments must be computed. The assumption is that subjects have no way of monitoring directly the presence of the solicited target in store and must infer it from a number of cues. Clearly, people can make predictions regarding many events on the basis of a variety of cues, and this ability need not imply the existence of an internal monitor where that information is directly stored. In a similar manner, the FOK may be based on an *inferential process*, con-

scious or unconscious, that uses a variety of cues to determine the likelihood that the solicited target is retained in memory and will be recognized or retrieved in the future (see, e.g., Costermans, Lories, & Ansary, 1992; Nelson & Narens, 1990). These cues may include familiarity with the general topic, retrieval of pertinent episodic information, and so on (see Nelson et al, 1984).

In the present article, I focus on one general cue for the FOK: the accessibility of pertinent information. According to the position advanced in this article, the cues for the FOK are to be found in the very information that is activated or accessed during the course of the search-and-retrieval process. Thus, the computation of FOK judgments is parasitic on the processes of retrieval. Whenever subjects interrogate their memory for a specific piece of information, a variety of clues come to mind. These include activations from the terms in the question; structural, contextual, and semantic attributes; fragments of the target; and so on (see Durso & Shore, 1991; Lovelace, 1987; Read & Bruce, 1982). These may serve to motivate and direct further search (see A. S. Brown, 1991). When a candidate answer is selected, such clues, the by-products of the retrieval process, can be used to evaluate the likelihood that that answer is correct (e.g., the ease of accessing the answer, the extent of corroborative evidence collected). However, even when retrieval fails, these very clues contain important information that can be used in judging whether the target can be recalled or recognized in the future. In general, FOK judgments are based on an attempt to extrapolate from the processes that occur during a retrieval episode to future retrieval episodes: If nothing comes to mind now, nothing will come to mind tomorrow. Even when subjects are asked to provide only FOK judgments about a difficult-to-remember target, they normally do so by attempting to search for the answer and observing the outcome. In such cases, FOK judgments would seem to be based on a sort of mini-simulation of the entire search-and-retrieval process. Naturally, because such judgments are often based on a fast sampling and inspection of momentary activations (see Morris, 1990), a biased sample may lead to an unwarranted FOK or to an unwarranted feeling of not knowing (see Koriati & Lieblich, 1977). This account of the mechanism underlying FOK is similar to the availability heuristic postulated by Tversky and Kahneman (1973) to explain how people estimate proportions or frequencies.

The cues that are used in computing FOK during a retrieval attempt can be classified into two types, those that have to do with the accessibility of information pertaining to the target and those that are based on the specific content of the information accessed. The term accessibility will be used here to subsume two major cues: the amount of information activated or accessed and its intensity (e.g., its ease of access, vividness, specificity, and persistence). It is proposed that in many situations, and particularly at the early stages of the search process, FOK judgments are primarily based on the mere accessibility of information. For example, in a TOT state, a strong FOK may ensue from the accessibility of partial fragments of the target and their persistent recurrence. In other cases, however, the subject may be able to evaluate the content of the information recovered, pitting different clues against each other and making de-

liberate, educated inferences about the plausibility that the solicited target will be subsequently recalled or recognized. This may sometimes be the case with real-world knowledge, particularly at the later stages of retrieval. Here the subject's inference may be more properly designated "judgment of knowing" rather than "feeling of knowing," and the underlying process may, in fact, be the same as that underlying the prediction of future events in general. The experimental work to be reported in the following paragraphs focuses specifically on the accessibility heuristic, which can operate even in situations where content-based considerations and deliberate plausibility judgments are likely to play a limited role.

The accessibility account of the FOK proposed here accords well with the theoretical position of Jacoby and his associates (Jacoby & Kelley, 1991; Jacoby, Kelley, & Dywan, 1989; Jacoby, Lindsay, & Toth, 1992; Kelley & Jacoby, 1990) of treating conscious experiences as constructions based on inferences and has much in common with the notion of a "fluency heuristic" that is assumed to underlie the experience of familiarity. According to this position, the subjective experience of remembering is not simply a product of a memory trace but instead relies on an inference. The cues for that inference are to be found in "aspects of one's own thoughts and behavior, such as the ease with which ideas come to mind" (Kelley & Jacoby, 1990, p. 49). Fluency of processing can serve as a valid cue for remembering, because prior experience makes current processing more fluent. However, fluent processing can also result from other sources, and when such sources are difficult to specify, fluent processing may lead to memory illusions (see Jacoby & Whitehouse, 1989; Whittlesea, Jacoby, & Girard, 1990). Thus, the work of Jacoby and his associates is a good demonstration of how the experience of remembering can be manipulated by altering current processing independently of past experience,

Metatheoretic Assumptions and Methodological Implications

Each of the two accounts of FOK delineated earlier is based on a different conceptual metaphor that carries distinct methodological implications. The metaphor underlying the internal-monitor view takes as its point of departure two basic features of the FOK phenomenon: first, the discrepancy between objective and subjective indexes of knowledge and, second, that, by and large, FOK judgments constitute valid indicators of actual knowledge. This pattern naturally suggests a specific dual-process metaphor of the FOK state: In this state, two independent memory processes are aimed at the same memory target, a memory-monitoring process (designated MEMO by Hart, 1966), attempting to detect its presence, and a search-and-retrieval process attempting to recall it. The FOK and TOT states represent memory blockage states (see Gruneberg, Smith, & Winflow, 1973; Jones & Langford, 1987), where the retrieval process is blocked while the memory-monitoring process remains intact.

When pushed to the extreme, the assumption of independent processes is taken to imply that the memory-monitoring process continues to tap the presence of the correct answer even when the answer retrieved is incorrect. For example, it has been

argued that the failure to retrieve the target in the TOT state stems, in part, from the interfering effect of "blockers" or "interlopers" that come to mind (Burke, MacKay, Worthley, & Wade, 1991; Jones, 1989; Reason & Lucas, 1984). The two-process metaphor implies that although such interlopers inhibit target recall, they do not necessarily detract from the subject's ability to monitor the presence of the correct target in the memory store.

In contrast, according to the accessibility account proposed here, subjects have no privileged access to an internal monitor that can inform them whether they know an answer that they are unable to recall. It is only by attempting to search for the solicited target that one can assess the likelihood that it is available in store and can be retrieved. Thus, FOK judgments must be computed on-line on the basis of clues accumulated during the initial stages of search and retrieval: The abortive attempt to retrieve the target leaves behind scattered debris that feed into a memory-monitoring process, which assesses the likelihood that the target will eventually be located. This process, then, is not independent of the retrieval process; if the latter goes astray, so will the former. This single-process view avoids the difficulties involved in the filing-cabinet conception of memory and is more compatible with the content-addressable view advocated by proponents of connectionist models of memory (e.g., McClelland & Rumelhart, 1985).

The two approaches sketched earlier have divergent methodological implications. Most important, from the point of view of the accessibility account, some of the previous FOK studies suffer from a fundamental methodological flaw: a confusion between the perspective of the subject and the perspective of the experimenter. Although this confusion reflects the pervasive influence of the dual-process metaphor, it also prevents a proper evaluation of the accessibility account. The confusion is manifested in two forms. The first applies to situations where subjects provide what appears to constitute a complete answer, whereas the second applies to situations where there is only partial recall of the intended target.

The Case of Complete Recall

Consider some of the procedural practices that follow from the recall-judgment-recognition (RJR) paradigm commonly used in the study of FOK (see Hart, 1965). Assume that a subject is presented with the question "What is the name of the tiny principality between France and Spain?" If the subject fails to provide any answer, then the experimenter solicits FOK judgments and tests for target recognition. However, if the response is "Andorra," then the experimenter will typically move on to the next item. Why is that? Presumably because of the assumption that if the experimenter knows that the answer is correct, the subject also knows that. In addition, of course, if the subject is right (i.e., if he or she produces the experimenter's intended target), then clearly there is no "memory blockage," and there is no sense in eliciting FOK judgments. However, as indicated earlier, a basic concern in the study of the FOK is whether subjects are accurate in monitoring their knowledge and, if so, what is the basis for their accuracy. Yet, in the procedure just described, a significant subset of the memory instances are elimi-

nated from further analysis merely because the experimenter knows that the subject is right.

In contrast, assume that the subject's response is "Lichtenstein." Now the experimenter's reaction will be completely different: He or she will typically proceed to solicit FOK judgments. This is because the subject's response is now seen to constitute a "retrieval failure," that is, a failure to recall the correct target. Thus, again the experimenter imposes his or her own perspective and interferes with the natural course of the monitoring process, selectively focusing on those responses that represent retrieval failures from the experimenter's point of view. This intervention may contaminate the subsequent FOK judgments and recognition performance, because the experimenter's behavior contains feedback about the correctness of the subject's answer (e.g., DaPolito, Guttenplan, & Steinitz, 1968; Krinsky & Nelson, 1985; Lupker et al., 1991; Schacter & Worling, 1985). Such contamination is all the more serious when subjects are forced to provide an answer to each item (e.g., Blake, 1973), as noted by Krinsky and Nelson (1985). For example, assume that a subject who has been asked to recall the capital of Israel is unable to decide between Tel Aviv and Jerusalem. When she finally opts for Tel Aviv, the experimenter informs her that the answer is incorrect and proceeds to solicit FOK judgments. Clearly, now she can be quite certain that she "knows" the correct answer (see Nelson & Narens, 1990).

These methodological problems reflect a basic conceptual dilemma whose solution must depend on one's pretheoretic stand. Thus, for those who subscribe to the dual-process metaphor, the unique feature of the FOK phenomenon lies in the discrepancy between the inability to retrieve a target and the ability to detect its presence in memory. Therefore, to demonstrate the phenomenon, one must first ascertain "retrieval failure," and this failure is implicitly defined as the inability to retrieve the correct answer. Consequently, it is a common practice in FOK studies to classify as retrieval failures both omission as well as commission errors, presumably because of the implicit assumption noted earlier that the memory-monitoring process can directly detect the availability of the target in store independent of the output of the retrieval process. Thus, it is assumed that even when a subject insists on an erroneous answer (e.g., that the capital of Australia is Sydney; see Gruneberg et al., 1973; Nelson et al., 1984), his or her FOK judgments nevertheless tap the presence of the correct target (i.e., the one intended by the experimenter).

This approach also implies that FOK must be validated against the subsequent recall, or recognition of the correct target. The alternative approach, that of using the subject's target as the criterion, is, in fact, the one adopted by R. Brown and McNeill (1966) in their classic study of the TOT phenomenon. In that study, the intended (experimenter's) target was revealed to the subjects at the end of each trial, and they indicated whether it was the one they were seeking or not. If not, they were asked to provide their actual ("effective") target if they could recall it. The partial information provided regarding a TOT target was then validated against the subject's target. Thus, the experimenter's target was never used as the criterion when the subject indicated that it was not the one he or she had been seeking. This solution has some limitations, because, as Koriat

and Lieblch (1974) showed, the accuracy of the partial information provided by subjects differed markedly when the subject's own target was the same as that intended by the experimenter than when it was not.

In FOK studies, too, there have been some attempts to take into account the subject's effective target in assessing FOK accuracy (see review by Krinsky & Nelson, 1985). Already in one of the early experiments, Hart (1966, Experiment 2) tried to avoid the dilemma mentioned earlier by defining retrieval failure operationally in terms of some mixture of subject-based and experimenter-based criteria. He forced subjects to guess the answer to each of the questions and then evaluated FOK's accuracy using only questions for which the answer was both incorrect (by the experimenter's standards) and also believed to be incorrect by the subject. Note, however, that forcing subjects to supply an answer (e.g., Blake, 1973; DaPolito et al., 1968; Hart, 1967a) again embodies the idea that the experimenter must first ensure that subjects really do not know the answer and only then try to demonstrate that they, in fact, do know that they know it.

These methodological policies would be untenable from the point of view of the accessibility account, where subjects' FOK judgments are assumed to be based on the very output of the retrieval attempt. According to this account, the more experimenters make sure that subjects have no memory of the target, the more experimenters deprive them of the very basis for their positive FOK: the output of the retrieval attempt. Thus, from the perspective of the subject, the demarcation line lies between "blanks" (omissions) and "nonblanks," rather than between "successful" (correct) and "unsuccessful" (omissions and commissions) retrievals. Therefore, it should not come as a surprise that FOK judgments are consistently higher following commission errors than following omission errors (Krinsky & Nelson, 1985).

The Case of Incomplete Recall

Consider next the case where subjects can provide only some partial information about the target. The hypothesis that FOK judgments rest on gaining access to partial information regarding the unrecalled target has been advanced by several authors (e.g., Blake, 1973; Eysenck, 1979; Koriat & Lieblch, 1977; Nelson, Leonesio, Shimamura, Landwehr, & Narens, 1982; Schacter, 1983; Schacter & Worling, 1985), but only a few studies tested this hypothesis in some systematic manner in connection with the RJR paradigm. Three such studies are Blake (1973), Eysenck (1979), and Schacter and Worling (1985), and their common methodological policy merits some attention. In Blake's study, FOK judgments were found to increase with the number of letters that the subject recalled about the target (a nonsense trigram). In Eysenck's study, subjects were asked to define rare words, and when unable to do so, their FOK judgments regarding the meaning of the word correlated with correct ratings of that word on the semantic differential. In the study of Schacter and Worling, subjects' FOK ratings regarding unrecalled words in a paired-associates task were correlated with correct ratings of the connotation of the word as good or bad.

Although all three studies confirmed the expected positive correlation between FOK and access to partial information, none of them represents a proper test of the accessibility hypothesis, because all involve the same methodological pitfall noted earlier of confusing the subject's and the experimenter's perspectives. Consider, for example, the procedure used by Blake (1973, Experiment 1): In each trial, subjects first memorized a three-letter trigram (e.g., LBN) and were later asked to report the letters. When recall was less than perfect, they provided FOK judgments and were then tested for recognition. The percentage of positive FOK judgments increased monotonically with the number of letters recalled, from 32%, when no letters were recalled, to 73%, when two letters were recalled. (Of course, FOK judgments were not solicited when all three letters were recalled, again because of the assumption discussed earlier.)

The problem with the methodology used by Blake (1973) is that subjects were forced to provide three letters ("so that overcautious withholding of correct information would be minimized" [p. 313]; see also Hart, 1966, 1967a), and the measure of partial information had to be the number of correct letters recalled. The scoring of the subject's report, of course, was done by the experimenter. This methodology implies a particular version of the partial information hypothesis: FOK judgments rest on the accuracy of the partial information accessible.

This implication, however, creates a dilemma: On the one hand, the force of the partial information hypothesis lies precisely in the assumption that subjects do not have direct access to the target's trace and therefore must infer target availability from partial recall. On the other hand, the experimental procedure used implicitly assumes that they do have direct access to the correctness of their partial recall. This makes the partial information hypothesis circular: Subjects feel that they know because they, in fact, know how much they know.

This problem also characterizes the studies of Eysenck (1979) and of Schacter and Worling (1985). In these studies too, partial knowledge was solicited through a forced-report procedure, so that "partial information" actually meant correct partial information. In fact, the finding of all three studies, that FOK judgments are correlated with correct partial information, has little explanatory power, because it is as puzzling as the finding that FOK judgments are positively correlated with correct recognition.

The accessibility account I propose in this article takes a more general form: It assumes that not only are subjects incapable of monitoring the availability of information in store but they are also incapable of monitoring directly the accuracy of the accessible information. Therefore, they often have to make do with the sheer amount of information accessible and with its intensity (e.g., ease of access). Thus, FOK judgments depend on the accessibility of information, as such, regardless of its correctness. To assess the amount of information accessible to the subject, a *free-report* procedure must be used, where subjects are allowed to volunteer only the partial information they can remember.

Explaining the Accuracy of the FOK The previous discussion leads to the distinction between two versions of the accessibility hypothesis. The first, which will be

designated the *target-retrievability hypothesis* (Schwartz & Metcalfe, 1992), assumes that FOK judgments are narrowly tuned to correct partial information about the target and, hence, that subjects have access to the accuracy of the recalled information. The second, is the *accessibility hypothesis*, which assumes that subjects cannot directly monitor the accuracy of the information that comes to mind and must base their FOK on the quantity and intensity of the information accessible.

How do these two hypotheses explain the accuracy of FOK judgments in predicting actual memory performance (Question 3 stated earlier)? Clearly, whereas the target-retrievability hypothesis may seem incomplete in explaining the basis for FOK judgments (Question 2), it does provide a straightforward explanation for their accuracy. This is because both correct partial recall and correct recognition are indicators of objective knowledge and ought to be positively correlated. In contrast, it is not immediately clear how the accuracy of FOK can be explained if FOK rests on the mere quantity of partial information accessible. I now address this question. (For ease of exposition, I ignore the contribution of intensity factors in the following discussion.)

In general, FOK judgments use the quantity, q , of information accessible at time $t1$ to predict the likelihood of correct memory performance, c , at time t . The accuracy of FOK judgments, that is, the correlation between q (at $t1$) and c (at t), should depend on two factors: first, how much of the information accessible at t is correct and, second, the extent to which accessibility at t is predictive of accessibility at $t1$.

Consider the first factor. Assume that subjects are presented with general-information or vocabulary questions, each requiring a one-word answer (Koriat & Lieblich, 1977, have used the term *memory pointer* to designate such questions as well as any cue that is intended to specify a particular memory entry), and are asked only to make a dichotomous FOK judgment whether or not they will recall the target at some later time (as in the Game Show paradigm of Reder, 1987). Assume further that subjects use a simple algorithm: Make a negative FOK if you can think of no reasonable answer and a positive FOK when some answer (any answer) comes to mind. The conditions for FOK's accuracy can now be specified: Given that some of the pointers elicit no answers (blanks), and assuming that subjects tend to provide at time t the same candidate answers as those available at time $t1$ (and draw the same blanks), then the accuracy of FOK judgments in predicting recall performance at time $t2$ should increase as a function of the ratio of correct to incorrect candidate responses at time $t1$. That is, the accuracy of FOK judgments is determined practically entirely by the correctness of the candidate answers. Note that what matters is not the overall difficulty of the test, that is, the proportion of items answered correctly out of the total number of items presented, but the proportion of correct answers out of the total number of answers accessed. In the terminology of Koriat and Goldsmith (1993a), these two indexes correspond to input-bound and output-bound measures of memory performance, respectively.

A typical result with most free-report memory tests is that correct responses represent a much larger proportion of the total number of responses reported than incorrect responses (see Koriat & Goldsmith, 1993a). This is no mere accident; it re-

flects the mundane fact that an item that has been stored in memory is more likely to give rise to correct than to incorrect (full or partial) reports. For example, in free-recall tasks, the number of correct responses generally exceeds the number of commission errors by a wide margin, and this is so even when subjects are encouraged to report all items that come to mind (see Bousfield & Rosner, 1970). Thus, free-recall responses, partial or complete, will be divided mostly between correct responses and blanks (omissions), with commission errors accounting for only a very small proportion of the answers. Under such conditions, a monitoring mechanism that relies solely on the accessibility of information, as such, is bound to be predictive of subsequent recall or recognition performance. Only under some contrived conditions, for example, those including a large number of "deceptive items," would subjects produce more incorrect than correct responses, and then they may also evidence an "unwarranted feeling of knowing" or endorse the incorrect response with great certainty (see Fischhoff, Slovic, & Lichtenstein, 1977; Koriat & Lieblich, 1977; Nelson et al., 1984).

A good example for the relationship between FOK's accuracy and the overall accuracy of the responses is provided by Koriat's (1975, 1976) studies. The results of the earlier study suggested that subject's FOK ratings monitor the correctness of subjects' responses in a task requiring the matching of antonyms in English with antonyms from noncognate languages. However, because in that study subjects performed better than chance, the possibility exists that FOK ratings actually monitored the "strength" of the match (e.g., its likelihood to be endorsed), not its accuracy. This possibility was supported in a subsequent study (Koriat, 1976), which was designed to deliberately include many items where subjects tended to prefer the incorrect translation. Whereas the correlation between FOK ratings and translation accuracy was positive for a subset of the items with predominantly correct responses, it was negative for a subset of the items where subjects' responses were predominantly incorrect. Thus, FOK accuracy depended critically on the overall accuracy of the matches.

The argument, then, is that in memory studies too, FOK judgments monitor the mere amount of information accessible. Nevertheless, they do differentiate between correct and incorrect responses simply because, by and large, accessible information is more likely to be correct than wrong. Therefore, the correlation between FOK judgments at time $t1$, and some measure of correct memory performance at time t , is more telling about memory performance than about metamemory performance.

Turning next to the second factor, it should be clear that the accuracy of FOK should be mitigated by any systematic differences between the information accumulated at time $t1$ and that which becomes available at time $t2$ (see also Morris, 1990). Such systematic differences can account for the impression that FOK judgments are based on a separate process from that underlying target retrieval. For example, a subject in a TOT state may retrieve partial information about the target and later realize that this information bears no relationship to the correct target that he or she eventually retrieves. Alternatively, a subject may initially express a negative FOK and later recall the correct answer spontaneously (e.g., the Don't Know-Got It-Correct

state, Koriat & Lieblich, 1974). Such changes over time need not imply that memory monitoring is independent of retrieval, because at every point in time FOK judgments monitor the accessibility of pertinent information at that point in time.

This idea can be illustrated with regard to the findings of Koriat and Lieblich (1977): When memory pointers (word definitions) were analyzed in terms of the memory states that they tended to precipitate, they were found to differ along two orthogonal factors—effectiveness in suggesting or eliciting the correct target ("objective knowledge") and degree of initial FOK ("subjective knowledge"). Taken at their face value, these results would seem to argue against the single-process view of the FOK. However, a close examination of the results suggests that the two factors may be interpreted as reflecting the distinction between the quantity, q , of information accessible at time t and the accuracy, c , of the information accessible at time t . Thus, for example, memory pointers eliciting accurate positive or negative FOK judgments typically provide the subject with an effective search plan that can allow him or her to zero in on the target or on the location in which it may reside. In terms of the present proposal, such pointers tend to induce selective focusing on correct as against incorrect partial information, thus contributing to FOK's accuracy. Other pointers, in contrast, tend to evoke a relatively large amount of incorrect clues at the initial stage of the search process (some originating from neighboring targets; see Jones, 1989; Reason & Lucas, 1984), resulting in an unwarranted FOK. Other pointers still elicit an unwarranted feeling of not knowing, apparently because they are characterized by a positively accelerating rate of information accumulation. Such changes in the time course of retrieval may explain why characteristics of the question (e.g., familiarity of the terms) may sometimes be more critical for initial FOK judgments than the reliability of the answer (Reder, 1987, 1988; Reder & Ritter, 1992; Schwartz & Metcalfe, 1992).

In summary, the accessibility account proposed here can explain both the basis for FOK judgments (Question 2) as well as the reason for their validity (Question 3). Furthermore, this account also permits addressing the question of FOK failure in predicting memory performance and allows specification of the conditions likely to produce high or low FOK accuracy.

An Accessibility Model of FOK and Some Predictions

I now sketch out a process model of the FOK on the basis of the ideas presented earlier. The model focuses on the effects of accessibility, disregarding factors that have to do with the content of the information accessible. I illustrate this with regard to the experimental procedure used in Experiment I, which is a modification of that used by Blake (1973; see also Hart, 1967a). In each trial, subjects memorized a four-letter string (e.g., TLBN), and after a short retention interval they were asked, to report the full target or as many letters as they could remember. Then they made FOK judgments, and their recognition memory for the target was tested.

This procedure conforms to the RJR paradigm introduced by Hart (1965), except for several notable modifications. First, unlike most previous studies, FOK judgments were always solicited regardless of the subject's performance on the initial re-

call test. (As noted earlier, eliciting FOK judgments only when the subject's answer is incorrect assumes that the subject has direct access to the correctness of his or her answer). Second, unlike the previous studies that tested the partial information hypothesis (Blake, 1973; Eysenck, 1979; Schacter & Worling, 1985), where partial knowledge was assessed through a forced-report procedure, here subjects were allowed the option to report as many letters as they could remember. Finally, the task chosen was one that minimizes both the effects of preexperimental subjective familiarity and the contribution of considerations having to do with the content of the information retrieved. In this manner, it was possible to focus only on the overall accessibility of information.

This experimental paradigm permits evaluation of the accessibility model of FOK depicted in Figure L. The model is presented in terms of a series of propositions, grouped into four categories: the determinants of FOK, the effects of memory strength, the predictive validity of FOK judgments, and the effects of ease of access. It should be noted that some of the predictions hold true only when certain conditions are satisfied, as is clarified in the General Discussion section.

The Determinants of FOK

The core assumption of the model is that FOK depends on the accessibility of partial information, regardless of its correctness. Accessibility is here defined to include two factors: amount of information and ease of access (as an example of an intensity cue). I deal first with the amount factor.

As far as the amount of partial information is concerned, two components should be distinguished, correct partial information (PI-C; e.g., the number of correct letters reported) and wrong partial information (PI-W; e.g., the number of wrong letters reported). The distinction between the two components is assumed not to be directly available to the subject, so that what matters is only PI-T, the total number of letters reproduced. It is proposed that

1. FOK increases with increasing PI-C (link *if*). It is essentially this prediction that has been evaluated and confirmed by Blake (1973), Eysenck (1979), and Schacter and Worling (1985).
2. However, a positive correlation is also predicted between




Figure J. The accessibility model of the feeling of knowing (FOK). (PI-C = correct partial information; PI-W = wrong partial information. The letters a-i represent links.)

FOK and PI-W (link *e*). This correlation is often masked by other contributions but should be observed when the effects of PI-C are controlled.

3. Successful access to partial information motivates further effort toward the retrieval of the complete target. This motivation may be similar to the motivation for task completion discussed by Lewin (1935; see Van Bergen, 1968). This effect may be mediated by enhancement of FOK (see Ryan, Petty, & Wenzlaff, 1982), as suggested by the finding that subjects spend more time searching for items they feel they know than for those they feel they do not know (Gruneberg, Monks, & Sykes, 1977; Lachman, Lachman, & Thronesbery, 1979; Nelson et al., 1984; Nelson & Narens, 1980).

4. Variables that increase PI-C generally reduce PI-W, so that PI-C and PI-W should generally correlate negatively across items and conditions (link *c*). This correlation, however, should vary depending on the origin of the erroneous partial information (see the General Discussion section).

Memory Strength as It Relates to Objective and Subjective Indexes of Knowledge

I use the term memory strength here in the sense that it is traditionally used in memory theories (see Wickelgren, 1970), without commitment to any specific theory of memory representation. Memory strength exerts a direct effect on actual memory performance and an indirect effect on FOK.

5. Memory strength affects likelihood of correct recognition (link *l*).

6. For targets stored in memory, PI-C will normally exceed PI-W. Furthermore, enhanced memory strength should generally increase PI-C (link *a*) and reduce PI-W (link *b*). More precisely, it should increase the ratio of PI-C to PI-W, thus improving the quality of the partial information accessible.

7. It follows that FOK judgments should generally increase with increases in memory strength. Previous work exploring this relationship has yielded somewhat inconsistent results, and the reasons are examined in the General Discussion section.

The Validity of FOK Judgments in Predicting Recognition Performance

8. PI-C and recognition should be correlated positively (link *l*) by virtue of the fact that they both depend on memory-strength (see Blake, 1973; Schacter & Worling, 1985). Therefore, the dependence of FOK on PI-C should be responsible for the success of FOK in predicting correct recognition.

9. In contrast, PI-W is responsible for FOK's inaccuracy. This is because PI-W is expected to be negatively correlated with recognition (link *g*) but positively correlated with FOK (link *e*). Thus, to the extent that FOK relies on PI-W, it should prove faulty, sometimes fostering unwarranted FOK (see Koriat & Lieblich, 1977).

10. Why then is the FOK-recognition correlation (link *h*) positive? This correlation reflects the degree of *resolution* (see Lichtenstein et al., 1982; Yaniv, Yates, & Smith, 1991), that is, the extent to which a subject can discriminate between items that are subsequently recognized and those that are not. This, in

turn, depends on the *variances* of the partial information components (see Nelson & Narens, 1990). Assuming that PI-C and recognition are positively correlated, then the correlation between FOK and recognition should be positive as long as variations in PI-C account for a greater proportion of the variance of PI-T. This will normally be the case for two reasons. First, by and large, a target that has been stored in memory is more likely to give rise to correct than to incorrect reports (see proposition 6). Therefore, PI-C is expected to also have a larger variance than PI-W. Second, in most situations, PI-C is more diagnostic of memory strength (link *a*) than is PI-W (link *b*). That is, interitem differences in memory strength are more systematically reflected in PI-C than in PI-W.

11. Factors that enhance memory strength (and overall memory performance) should also increase the accuracy of FOK in predicting subsequent recognition performance (see, e.g., Carroll & Nelson, 1993). Such factors should increase the ratio of PI-C to PI-W, and if this increase also leads to a greater contribution of PI-C to the PI-T variance, then improvement in FOK accuracy should be expected.

Ease of Access of Partial Information

FOK is also affected by the intensity of the accessed information, that is, the ease with which partial clues come to mind, the subjective vividness and specificity (e.g., degree of detail) of the information, and its persistent, spontaneous recurrence.

The contribution of intensity factors can help explain why subjects are not only accurate in predicting recognition performance (prospective monitoring) but also in judging the correctness of their complete or partial reports (retrospective monitoring). They do so, in part, on the basis of the intensity of the accessed information. Here I focus only on the ease with which information comes to mind (e.g., retrieval latency). Most of the assumptions concerning ease of access parallel those pertaining to the amount of partial information.

12. Correct partial or complete information is retrieved with greater ease than incorrect information, so that ease of access is diagnostic of the accuracy of the accessed information.

13. Ease of access of partial information affects FOK independent of the amount of information retrieved (see Costermans et al., 1992; Nelson et al., 1982; Nelson & Narens, 1980).

14. It follows that a correct piece of partial information should make a greater contribution to FOK than an incorrect piece, and therefore FOK judgments ought to be diagnostic of the accuracy of the information retrieved.

15. Manipulations that enhance overall memory performance also increase the ease of accessing information. As a result, a positive correlation may be expected between the amount of information retrieved and the ease with which it is retrieved.

16. Both recognition performance and ease of access of partial information should vary with memory strength. Therefore, FOK's reliance on ease of access should improve the accuracy of FOK in predicting recognition performance.

17. Rate of information accrual is critical for both FOK and its accuracy. When a memory pointer immediately evokes a wealth of associations, but the association reserves are quickly depleted, it may engender an unwarranted positive FOK. The

reverse situation, where information accumulates at a negatively accelerating rate, would often lead to an unwarranted feeling of not knowing (see Koriat & Lieblich, 1977; Morris, 1990).

This model was tested in three experiments that together addressed most of the propositions listed earlier.

Experiment 1

The experiment was modeled after the early experiments by Blake (1973) and Hart (1967a). In each trial, subjects memorized a four-letter nonsense string and were then asked to recall as many letters as they could from it. Finally, they provided FOK judgments, and their recognition memory was tested. This procedure differs from that of Blake in two respects: First, recall was tested in a free-report format, and, second, FOK judgments were always solicited regardless of the accuracy of the letters reported.

Method

Subjects. Thirty Hebrew-speaking University of Haifa undergraduate students participated in the experiment; 19 received course credit and 11 were paid.

Materials. Targets were 40 four-place English consonant strings (tetragrams). They were generated randomly except that each of the 20 consonants appeared in at least 7 and at most 9 of the tetragrams, and none of the letters was repeated within the same tetragram. Letter strings likely to evoke specific associations in English or Hebrew were avoided. For the distractor task, Stroop-type items (Stroop, 1935) were used, with four colors and their Hebrew color names. The recognition test included the target and seven lures, which were constructed as follows: One lure (the contrast lure) was selected that shared no letters with the target. The remaining six lures were composed of the eight letters of the target and the contrast lure, so that the ratio of target letters to contrast letters was 1:3 in two lures, 2:2 in two lures, and 3:1 in two lures. This construction was used so that subjects would not be able to guess the correct target from the relative frequency of different letters among the distractors. The order of the letters within each distractor (including the correct alternative) was random except that no target letter occupied the same location as in the original study target. For example, for the target FKDR, the distractors were RDFK, RFSC, BCSF, SBJC, DBFK, KjBC, RDJS, and RJKD.

Procedure. The experiment was conducted on an Apollo DN-4000 workstation. Each trial included four Stroop items (to reduce interference from the preceding target string), followed by the target string, followed by a series of Stroop items for 18 s. All stimuli appeared at the center of the screen. Subjects named the color ink as fast as they could into a voice-activated microphone, and the item was replaced by the next stimulus 500 ms after the response. The target string to be memorized appeared for 1,000 ms, with the letters arranged vertically. (In exploratory work, the vertical arrangement had been found to be more conducive to partial report than the horizontal arrangement.) After 600 ms from the offset of the target, the presentation of the Stroop items continued until 18 s elapsed.

The instruction "Recall the target" appeared then on the screen, with four squares underneath, vertically arranged. Subjects were instructed that on each trial they could gain 1 point for each correct letter but would win nothing on that trial if they report even a single wrong letter. (Without these special instructions, subjects in the exploratory work, presumably because of testing habits, had tended to either report four letters or no letters at all.) It was indicated that the order of the letters

was immaterial. The reported letters were displayed inside the four squares as they were typed in by the experimenter.

Subjects then indicated their FOK on a 0-100 scale, expressing the chances of identifying the target. A detailed description of the structure of the recognition test was given. For the recognition test, the eight distractors appeared on the screen, randomly ordered, each displayed in a vertical format.

The experiment included 2 practice and 40 experimental trials. The order of the target items was randomly determined for each subject, with the restriction that successive targets would not contain any letters in common. There was a short break after the first 20 trials.

Results and Discussion

The results are analyzed roughly according to the three questions posed in the introductory paragraphs: the predictive validity of FOK, the process underlying FOK judgments, and the basis for FOK accuracy.

The predictive validity of FOK judgments. The relationship between FOK and recognition (link, h in Figure 1) was first evaluated by pooling data across all subjects and items ($N = 1,200$), using the procedure applied in calibration studies (see Lichtenstein et al., 1982). Items were grouped into 10 FOK categories (1-10, 11-20, and so on), and Figure 2 presents the percentage of correct recognition for these categories. Calibration was very good, without the overconfidence bias typically found in studies of subjective confidence. The overall correlation ($N = 10$) between FOK and percentage correct was .97 ($p < .0001$). This correlation remained high, .90 ($N = 10$, $p < .001$), when responses with PE-C = 4 were excluded from the analysis (as is the common practice in FOK studies). Note that some of the

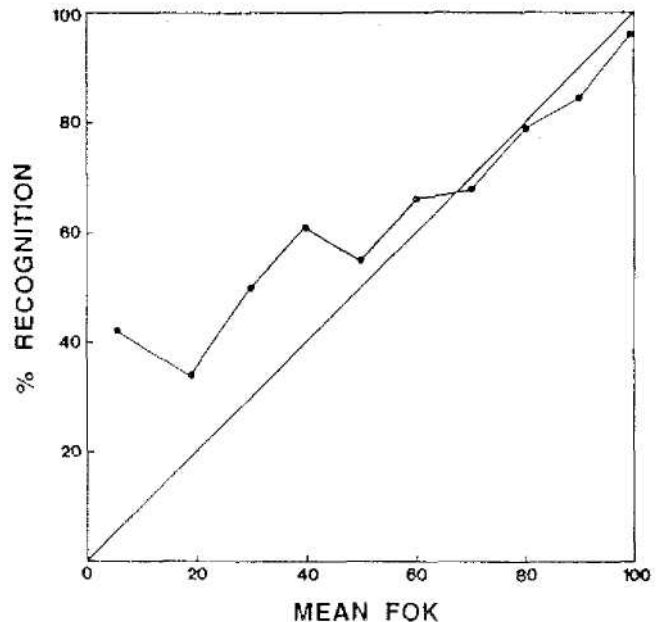


Figure 2. Calibration curve relating the percentage of correct recognition to feeling-of-knowing (FOK) judgments across subjects and items (Experiment 1).

reported FOKs were below 10%, although the likelihood of choosing the correct alternative by chance was .125.

Each point in Figure 2 is based on a different combination of subjects and items. Because of the very large intersubject variability in both mean FOK judgments (range = 35.5-98.2) and percentage correct recognition (range = 35%-97.5%), it was important to examine the intraindividual correlation between FOK and correct recognition. Each subject's FOK judgments were divided at the median, and the gamma measure recommended by Nelson (1984) was calculated for each subject. This correlation averaged .55, significantly different from zero, $t(29) = 9.06, p < .0001$. Gamma was negative for 1 subject, zero for 1 subject, and positive for the remaining 28 subjects.

The relationship between FOK and the amount of accessible information. For each subject, all trials were classified according to the number of correct letters reported (PI-C) and the number of wrong letters reported (PI-W). Table 1 presents the number of subjects yielding each of the patterns as well as the total number of such patterns across all subjects and trials. The results presented in this table illustrate the "fragmentary data problem" characteristic of studies on partial recall (see A. S. Brown, 1991; R. Brown & McNeill, 1966). Therefore, the analyses of these data require some departure from the conventional methods.

Mean FOK judgments were calculated for each subject for each of the response patterns demonstrated by him or her, and these means were averaged across all subjects evidencing each of the patterns. The group means are plotted in Figure 3 in two different formats, once (top panel) as a function of PI-C, with PI-W as a parameter, and once (bottom panel) as a function of PI-W, with PI-C as a parameter.

Consider first the effect of PI-C. FOK judgments increased systematically with increasing PI-C (link *d* in Figure 1). Pair comparisons of all mean FOK judgments for cells sharing the same PI-W value indicated that for all 20 possible comparisons the FOK means were higher for responses with higher PI-C values ($p < .0001$ by a binomial test). Interestingly, the high FOK mean (90.0) associated with the 4-0 pattern (i.e., four incorrect letters) was based on two observations, each for a different subject, in both of which the string reported was identical (i.e.,

same letters and in the same order) to the target presented on the preceding trial (see the General Discussion section).

Because it is a common practice not to solicit FOK judgments when the subject's answer is correct, the aforementioned analysis was repeated excluding trials with PI-C = 4. (Note that these trials were not associated uniformly with FOK judgments of 100.) The 16 pair comparisons thus remaining were all in the expected direction ($p < .0001$ by a binomial test).

Turning next to the relationship with PI-W (link *e* in Figure 1), it can be seen (Figure 3, bottom panel) that FOK judgments tend to increase with increasing number of letters when PI-C is held constant. That is, the more incorrect letters reported, the stronger the FOK. Thus, of the 20 pair comparisons between cells sharing the same value of PI-C, 19 were in the expected direction ($p < .0001$ by a binomial test). (The results remain unchanged when PI-C = 4 is excluded.)

It is important to stress that the overall correlation between FOK judgments and the number of incorrect responses is generally negative (when PI-C is not controlled): Across all subjects and items, mean FOK judgments for items associated with PI-W values of 0, 1, 2, 3, and 4 were 68.1, 53.9, 47.4, 39.4, and 90.0, respectively. Thus, except for the last value, which is based on the two observations mentioned earlier, mean FOK judgments decreased with increasing number of incorrect responses (see the following paragraphs).

In summary, FOK judgments increase with the amount of information accessible regardless of the accuracy of that information.

Accessibility of information as a predictor of recognition performance. To explain the validity of FOK in predicting recognition, it is necessary to examine the correlations between recognition scores, on the one hand, and PI-C and PI-W, on the other (links/and *g*-in Figure 1).

The analyses reported earlier for FOK judgments (see Figure 3) were repeated using percentage of correct recognition as the dependent variable. Mean recognition performance is plotted in Figure 4. It can be seen (top panel) that correct recognition increases with increasing PI-C. Pair comparisons over all points sharing the same PI-W value indicated such increase for 17 out of the 20 comparisons ($p < .001$ by a binomial test).

Table 1

Number of Subjects Yielding Each of the Combinations of PI-C and PI-W and the Total Number of Such Patterns Across All Subjects and Trials (Experiment 1)

PI-W	PI-C									
	0		1		2		3		4	
	<i>n</i>	Total no. of such patterns	<i>n</i>	Total no. of such patterns	<i>n</i>	Total no. of such patterns	<i>n</i>	Total no. of such patterns	<i>n</i>	Total no. of such patterns
0	10	23	21	73	26	187	30	222	29	380
1	10	18	25	63	24	103	23	61		
2	8	12	13	33	8	15				
3	4	5	2	3						
4	2	2								

Note. PI-C = correct partial information; PI-W = wrong partial information.

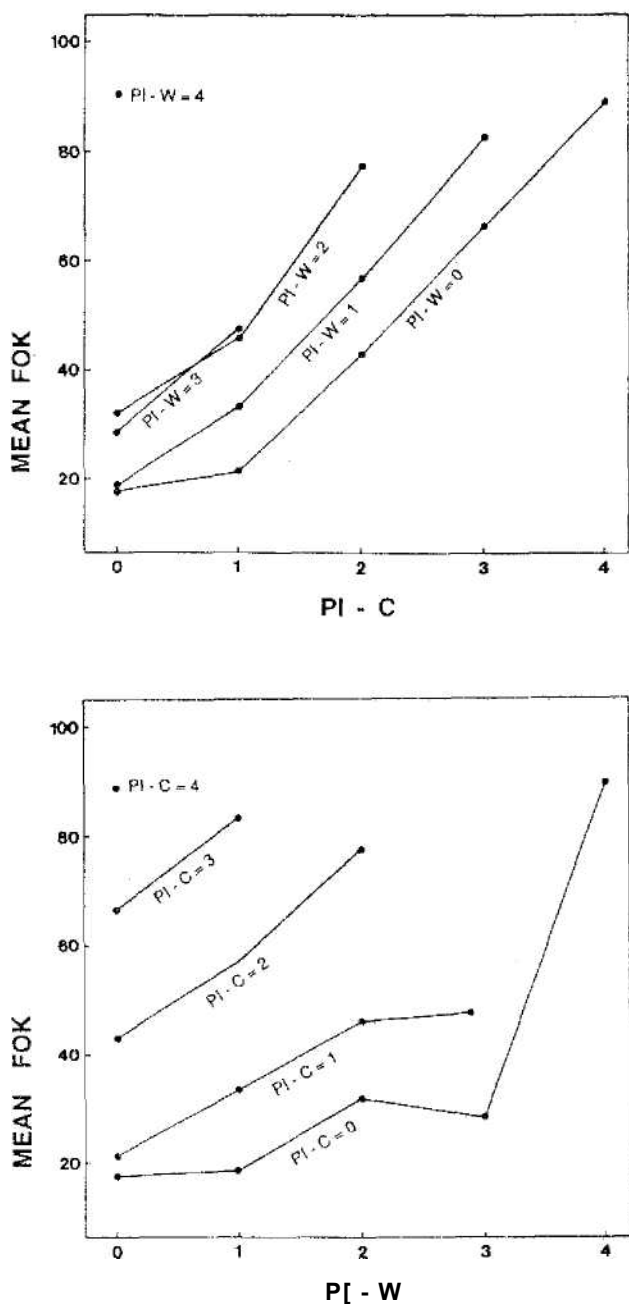


Figure 3. Mean feeling-of-knowing (FOK) judgments plotted (top panel) as a function of correct partial information (PI-C), with wrong partial information (PI-W) as a parameter, and (bottom panel) as a function of PI-W, with PI-C as a parameter (Experiment 1).

In contrast, the opposite trend is suggested for PI-W (bottom panel), with correct recognition decreasing with increasing number of letters reported. Although this relationship was somewhat less systematic than that obtained with PI-C, pair comparisons over all points sharing the same PI-C values indicated such a decrease in 15 out of 19 pairs (there was one tie; $p < .01$ by a binomial test).

The accuracy and inaccuracy components of FOK. The following analyses were carried out to estimate the overall correlations between the two components of partial information, on the one hand, and FOK and recognition, on the other hand.




Figure 4. Mean correct recognition plotted (top panel) as a function of correct partial information (PI-C), with wrong partial information (PI-W) as a parameter and (bottom panel) as a function of PI-W, with PI-C as a parameter (Experiment 1).

Consider first the relationship between FOK and partial information. To estimate the overall relationship between FOK judgments and each of the two components of partial information (PI-C and PI-W) when controlling for the effects of the other component, the following procedure was used. The FOK means presented in Figure 3 were first standardized for all points having the same value on the controlled component, and then an overall correlation was calculated between the other component and FOK. Thus, to examine the relationship between FOK and PI-C, the FOK means for each set of points with the same PI-W value were first transformed into standard scores with $M = 100$ and $SD = 10$. Figure 5 (top panel) represents the regression line relating the standardized FOK means to PI-C. The overall correlation across all 15 points (closed circles) was $.83$ ($0 < .0001$). This correlation remained high ($.87$, $N = 14$, $p < .0001$) when the analysis was repeated, excluding trials with $PI-C = 4$.

A similar analysis was carried out to evaluate the relationship of FOK with PI-W, by first standardizing the means of all points with the same PI-C value (Figure 3, bottom panel). The new means are also plotted in the top panel of Figure 5 (open circles) as a function of PI-W. The overall correlation across the 15 points was positive and quite high: $.76$ ($p < .002$). The respective correlation after elimination of trials with $PI-C = 4$ was $.79$ ($N = 14$, $p < .001$). Thus, the correlations between FOK and the amount of accessible information are positive regardless of the accuracy of the information retrieved.

Turning next to the relationship between partial information and recognition, similar analyses to those reported above for FOK, were carried out for recognition, and their results are displayed in the bottom panel of Figure 5. Thus, to examine the relationship between PI-C and recognition, the means presented in Figure 4 were first standardized for all points with the same PI-W value. The means of these standard scores are plotted in the bottom panel of Figure 5 (closed circles) as a function of PI-C, where the increase in recognition with increasing PI-C can be readily observed. The correlation ($JV = 15$) is $.61$ ($p < .02$).

The relationship between PI-W and correct recognition was analyzed in a similar manner by first standardizing the means of each set of points with the same PI-C value. As may be seen in Figure 5 (bottom panel, open circles), the correlation now is negative: $.52$ ($N = 15$, $r < .05$).

In summary, a comparison of the results presented in the top and bottom panels of Figure 5 tells the entire story: PI-C is related in the same way to FOK judgments and correct recognition; PI-W, in contrast, is positively related to FOK and negatively related to recognition. Thus, the dependence of FOK on PI-C is responsible for its success in predicting correct recognition, whereas its dependence on PI-W is responsible for its inaccuracy (propositions 8 and 9 of the model).

The mediating role of overall amount of accessible information. If FOK judgments depend on both PI-C and PI-W, why does PI-T, the overall amount of information recalled, nevertheless serve as a valid predictor of recognition performance? I first examine the relationship between PI-T and FOK and then show that PI-T itself is predictive of recognition. Finally, I offer

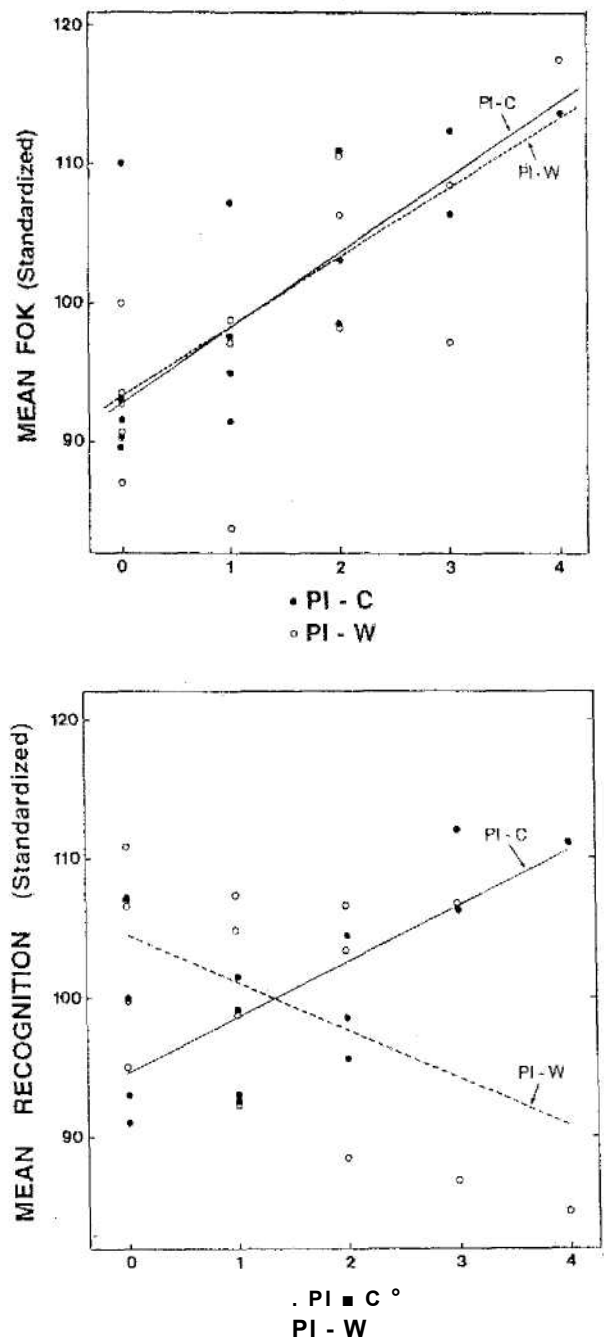


Figure 5, Regression lines relating the standardized feeling-of-knowing (FOK) means (top panel) and the standardized recognition memory means (bottom panel) to correct partial information (PI-C; closed circles) and wrong partial information (PI-W; open circles; Experiment 1).

an explanation of this pattern in terms of the relative contributions of PI-C and PI-W,

Consider first the relationship between FOK and PI-T. For each subject, mean FOK judgments for items with different PI-

T values were calculated, and these were averaged across all subjects for whom the respective means were available. The overall FOK means for PI-T values of 0, 1, 2, 3, and 4 were (number of subjects in parentheses) 17.9 ($n = 10$), 20.8 ($n = 21$), 40.1 ($n = 26$), 63.7 ($n = 30$), and 87.6 ($n = 30$), respectively. Thus, FOK judgments increase systematically and strongly as a function of the mere number of letters reported.

Next, consider the relationship between PI-T and recognition. Mean recognition memory for items associated with PI-T values of 0, 1, 2, 3, and 4 were 11.6, 53.2, 51.2, 69.7, and 90.0, respectively. Although the rank-order correlation is not perfect, it is clear that the mere number of letters retrieved is a very powerful predictor of recognition memory. Thus, although one component of PI-T is correlated positively with recognition, whereas the other is correlated negatively, the overall correlation between PI-T and recognition is clearly positive. Why is that?

The answer lies in the relative contributions of PI-C and PI-W to PI-T. These contributions can be estimated simply from the correlations between PI-T and each of its two components. The results disclosed a predominant contribution of PI-C: The mean within-subject correlation between PI-T and PI-C was very high, .83, $t(29) = 39.64$ for the difference from 0 ($p < .0001$). In contrast, the mean correlation between PI-T and PI-W was .00 (!). Note that PI-C and PI-W were negatively correlated: $-.51$, $t(29) = 13.37$, $p < .0001$.

It is clear, then, that the predictive validity of PI-T is almost entirely due to PI-C. There are two reasons for this state of affairs. First, items committed to memory are more likely to give rise to correct than to incorrect reports (proposition 6 of the model). This was indeed the case in Experiment 1, where 88.8% of all reported letters were correct. This is apparently responsible for the fact that the average variance of PI-C was about three times as large (1.08) as that of PI-W (0.33). Second, differences in memory strength are more likely to contribute systematically to PI-C than to PI-W. In fact, the high negative correlation between PI-W and PI-C suggests that much of the systematic variance of PI-W is due to its correlation with PI-C (see the General Discussion section).

In conclusion, the total number of letters accessible serves as a good predictor of recognition because most of the variance of PI-T is due to correct accessible information. This is illustrated in Figure 6, which compares the amount and type of partial information associated with successful recognition. It may be seen that successfully recognized targets are associated with the report of a larger number of correct letters, $f(29) = 12.09$, $p < .0001$, and a smaller number of incorrect letters, $f(29) = 6.72$, $p < .0001$, than unrecognized targets. However, because PI-C represents a larger portion of PI-T, the latter turns out to also increase positively with correct recognition, $t(29) = 8.45$, $p < .0001$.

Figure 6 also suggests that not only does PI-T increase with increasing memory strength but its composition also changes toward a greater representation of PI-C: The proportion of correct letters among those reported was 93% for recognized targets, compared with 73% for unrecognized targets.

Individual differences. Another approach to the issue discussed in the previous section is to compare performance of subjects with good and poor memory performance. The as-




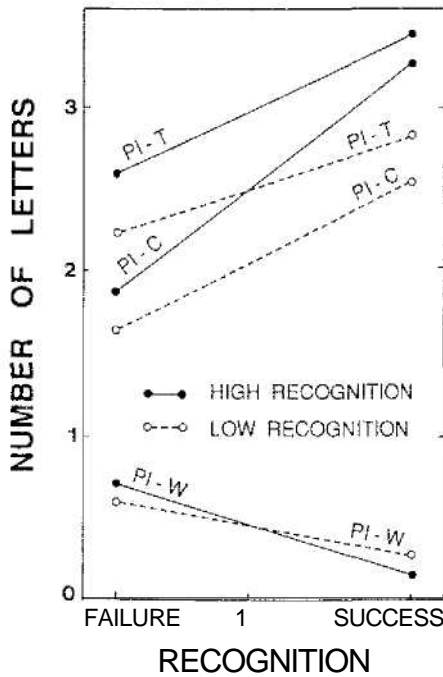
Figure 6. Mean correct partial information (PI-C), wrong partial information (PI-W), and total number of letters reproduced (PI-T) for recognized and not recognized targets (Experiment 1).

sumption is that average memory strength is higher for the former than for the latter subjects. Thus, we may ask, as Lichtenstein and Fischhoff (1977) did: Do those, who know more also know more about how much they know?

Subjects were divided into two groups according to their overall recognition performance. The low-recognition group included 14 subjects with 70.0% or less correct performance, while the high-recognition group included the remaining 16 subjects. Mean FOK was higher for the high (64.5) than for the low group (53.2). In addition, however, these judgments appeared to have a higher predictive validity in the high than in the low group. Thus, mean gamma correlation was .40 for the low-recognition and .67 for the high-recognition subjects, $t(28) = 2.39$, $p < .05$.

The results presented in Figure 7 explain the higher FOK accuracy of high-recognition subjects. This figure is constructed in a similar manner to Figure 6, except that it contrasts the partial recall performance of the low- and high- recognition groups. It may be seen that PI-C is higher for recognized than for unrecognized targets, $F(1, 28) = 172.63$, $p < .0001$, but this effect is stronger for the high than for the low group, $F(1, 28) = 7.26$, $p < .02$, for the interaction. In parallel, mean PI-W was lower for recognized than for unrecognized targets, $F(1, 28) = 46.93$, $p < .0001$, but again this effect tended to be stronger for the high group, $F(1, 28) = 3.29$, $p < .10$, for the interaction.

Another way to express this difference is to examine the mean FOK judgments associated with targets that were ultimately recognized and those that were not. Mean FOK judgments for



7. Mean correct partial information (PI-C), wrong partial information (PI-W), and total number of letters reproduced (PH) for recognized and not recognized targets, for low-recognition and high-recognition subjects (Experiment 1).

the two types of targets were 46.8 and 59.5, respectively, for the low group and 52.1 and 76.8, respectively, for the high group. A two-way analysis of variance (ANOVA) on these means yielded $F(1, 28) = 11.19, p < .005$, for the interaction, indicating that FOK judgments were more efficient in the high than in the low group in distinguishing between targets that were ultimately recognized and those that were not.

These results suggest that improved memory is associated with improved metamemory and that this association is mediated by the improved quality of the accessible partial information (proposition 11 of the model),

Monitoring the accuracy of partial information. It was proposed that ease of retrieval may add to the predictive validity of FOK by serving as a cue for the accuracy of the accessed information. This implies that FOK judgments should also be sensitive to the accuracy of partial information, not only to its amount {proposition 14}.

Two analyses confirm this prediction. First, in Figure 8, mean FOK judgments are now plotted as a function of PI-C with PI-T held constant. It may be seen that except for one deviant point (that is based on the two discrepant observations mentioned earlier), FOK judgments increase systematically with the number (or proportion) of correct letters. This increase is evident for 16 out of the 20 pair comparisons ($p < .01$ by a binomial test). Thus, a correct letter makes a stronger contribution to FOK than an incorrect letter.

A second analysis supported the same conclusion. When examination was confined to trials for which subjects provided

four letters, then FOK judgments were significantly higher when all letters were correct ($M = 88.6$) than when some of the letters were incorrect ($M = 82.0$), $t(23) = 4.72, p < .0001$ (using only subjects for whom both means were available). The same was true for PI-T = 3 ($M = 64.8$ and 55.1 , respectively), $t(23) = 3.04, p < .01$, and for PI-T = 2, ($M = 42.7$ and 34.1 , respectively), $f(24) = 3.16, p < .005$, but not for PI-T = 1 ($M = 21.5$ and 18.8 , respectively), $f(20) = 0.94, ns$. An overall two-way ANOVA, PI-T X Recall Accuracy (perfect recall vs. not perfect) for unequal n s, yielded $F(3, 29) = 226.98, p < .0001$ for PI-T, $i^2(1, 29) = 21.93, p < .0001$ for recall accuracy, and $F < 1$ for the interaction,

In conclusion, the results of Experiment 1 are consistent with the postulates of the accessibility model concerning partial recall. It was shown that FOK judgments depend on the amount of information retrieved about the target regardless of its accuracy but that, nevertheless, these judgments are successful in predicting recognition performance. This success seems to ensue primarily from the fact that the retrieved information is more likely to be correct than to be wrong. In summary, then, the results indicate that there is no need to invoke the notion of direct monitoring of the memory trace to explain the validity of FOK judgments.

Despite the fact that the results were clearly in agreement with the model, it is important to treat them with caution because of some of the limitations of the experimental procedures used in the present study. First, the results of Experiment 1, as




Figure 8. Mean feeling-of-knowing (FOK) judgments as a function of correct partial information (PI-C), with the total number of letters reproduced (PI-T) as a parameter (Experiment 1).

well as most of those of the following experiments, are correlational in nature. Therefore, it is important to extend investigation to other experimental paradigms that allow the testing of some of the model's predictions through experimental manipulation. Second, the task of Experiment 1 (see also Experiment 2) was somewhat more artificial than some of the other tasks commonly used in the study of the FOK. This task was found to be particularly suited to the aims of the present study because it afforded assessment of all of the critical variables on each trial and, in particular, the amount of partial information accessed. In addition, it minimized the contribution of content factors and preexperimental familiarity. In this manner, it was possible to address most of the model's predictions within a single design. However, it is important to establish the utility of the accessibility model with less constraining procedures and with tasks that have a greater ecological validity. Note, though, that many of the results already reported with other tasks {e.g., answering general-information questions, see Nelson & Narens, 1990} do appear to be quite consistent with the present model.

Experiment 2

Experiment 2 was designed to evaluate the predictions of the accessibility model concerning the effects of memory strength and ease of access. The experiment was similar to Experiment 1 with three modifications. First, five-place strings were used to produce a wider range of partial recall and reduce the possibility of a ceiling effect on PI-W. Second, retention interval was manipulated to test the effects of memory strength. Third, latency of recall initiation was recorded.

Memory strength is expected to increase the ratio of P.I-C to P.I-W, thereby improving FOK's accuracy (propositions 6 and 11 of the model). Ease of access is expected to be diagnostic of the accuracy of the accessed information (proposition 12), to make an independent contribution to FOK judgments (proposition 13), and thus to allow subjects to monitor the accuracy of partial knowledge (proposition 14). It is also expected that ease of access of correct partial information should be affected by retention interval (proposition 15) and that, reliance on ease of access should contribute to FOK accuracy (proposition 16).

Method

Subjects. Twenty-four University of Haifa undergraduate students participated in the experiment for course credit. None had participated in the previous experiment.

Materials. Targets were 40 five-place English consonant strings, derived by adding one consonant letter to each of the tetragrams used in Experiment 1. The added letter was chosen randomly from the letters not included in the tetragram so that a different letter was added to two targets. Letter sequences that evoke specific associations were avoided. The recognition test consisted of 10 five-letter alternatives, and the lures were constructed according to the same scheme as in Experiment 1.

Procedure. The apparatus was the same as in Experiment 1, and the procedure was similar except for the following: First, the retention (Stroop) interval was 15 s for half of the items and 25 s for the remaining items, with a random assignment of intervals to items for each subject. Presentation order was random, except that it was the same for each pair of subjects, and that for that pair, retention interval was counterbalanced for each item.

Second, latency to initiate recall {i.e., pronounce the first letter} was measured by having subjects speak their responses into a voice-operated microphone. When the subject announced recall termination, by saying "that's it," the experimenter pressed a key, and the subject then reported his or her FOK judgment.

Results and Discussion

I first briefly summarize those aspects of the results that constitute a replication of those of Experiment 1.

The predictive validity of FOK judgments. The calibration functions were practically identical for the two retention intervals (Figure 9) and were similar to the function obtained in Experiment 1 (Figure 2). The FOK-recognition correlations for the grouped data were .82 ($p < .005$) and .93 ($p < .0001$) for the short and long intervals, respectively. The somewhat higher correlation for the longer interval is consistent with Nelson and Dunlosky's (1991) finding that judgments of learning are more accurate when made shortly after study than when made immediately after study.

The within-subject gamma correlation (across both retention intervals) between (dichotomized) FOK and recognition averaged .47, significantly different from zero, $t(23) = 1.13$, $p < .0001$. It was negative for 2 subjects and positive for the remaining 22 subjects.

The relationship between FOK and the amount of partial information accessible. Table 2 presents mean FOK judgments for each combination of PI-C and PI-W. FOK increased with increasing PI-C: Of the 34 pair comparisons between cells sharing the same PI-W values, 32 were in the expected direction ($p < .0001$ by a binomial test). There was also a tendency for FOK

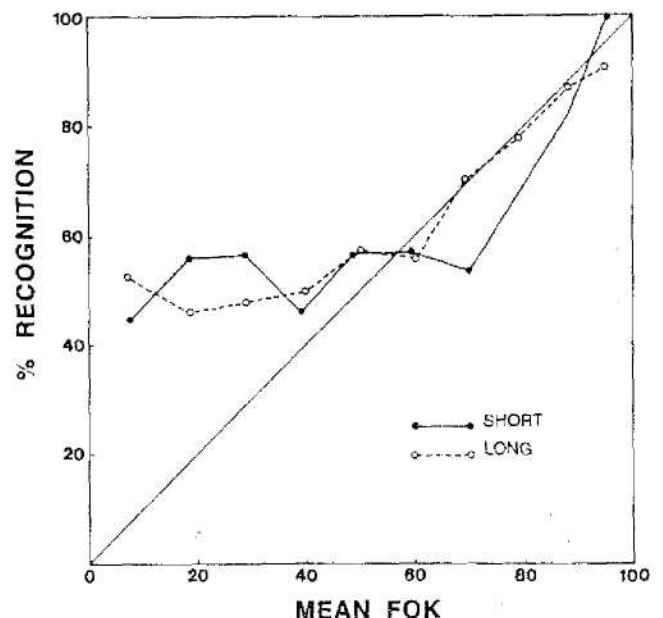


Figure 9. Calibration curve relating the percentage of correct recognition to feeling-of-knowing (FOK) judgments across subjects and items for the short and long retention intervals (Experiment 2).

Table 2
*Mean Feeling-of-Knowing Judgments as a Function of the Number of Correct Letters (PI-C)
 and the Number of Incorrect Letters (PI-W) Reported (Experiment 2)*

PI-W	PI-C											
	0		1		2		3		4		5	
	<i>M</i>	<i>n</i>	<i>M</i>	<i>n</i>	<i>M</i>	<i>n</i>	<i>M</i>	<i>n</i>	<i>M</i>	<i>n</i>	<i>M</i>	<i>n</i>
0	6.1	11	14.6	15	28.4	24	46.1	24	61.6	21	51.2	16
1	8.6	8	26.4	21	37.5	21	54.7	16	40.8	11		
2	24.0	5	36.0	15	56.5	8	84.2	3				
3	20.0	1	70.0	1	83.3	3						
4	—	0										
5	—	0										

to increase with increasing PI-W when PI-C was held constant. Thus, of the 22 pair comparisons between cells sharing the same value of PI-C, 20 were in the expected direction ($p < .0001$ by a binomial test).

To compare the effects of PI-C and PI-W when controlling for the other component, the standardization procedure used in Experiment 1 (see Figure 5) was applied. The overall correlation between PI-C and FOK, with PI-W controlled, was .84 ($JV = 18, />< .0001$). This correlation remained high (.89, $N = 17, p < .0001$) when the analysis was repeated excluding trials with PI-C = 5. The corresponding correlation between PI-W and FOK, with PI-C held constant, was .75 ($N = 18, p < .0005$). The respective correlation after elimination of trials with PI-C = 5 was .78 ($N = 17, p < .0005$).

In summary, the results indicate that FOK judgments increase with the amount of information accessible, irrespective of the accuracy of this information.

Partial information as a predictor of recognition memory performance. It can be seen (Figure 10, top panel) that recognition memory increased with increasing PI-C. Pair comparisons over all points sharing the same PI-W value indicated such an increase for 33 out of the 34 comparisons (there was one tie; $p < .0001$ by a binomial test).

In contrast, recognition memory evidenced a slight tendency to correlate negatively with PI-W (Figure 10, bottom panel). Pair comparisons over all points sharing the same PI-C values indicated a decrease in 14 out of the 24 pairs (there was one tie). The difference was not significant by a binomial test.

When the recognition means depicted in the top panel of Figure 10 were first standardized over all points with the same PI-W value, the correlation between recognition means and PI-C was .89 ($iV = 18, p < .0001$). When the analysis was carried out separately for the short and long retention intervals, the corresponding correlations were .91 and .87, respectively. In contrast, the correlation between recognition and PI-W when PI-C was controlled was $-.03$ ($\ll .?$). This correlation, calculated separately for the short and long retention intervals, was $-.18$ and $-.05$, respectively. Thus, unlike the results of Experiment 1, where there was a tendency for recognition to decrease with increasing PI-W, here this tendency was much less clear.

In summary, PI-C correlated positively with both FOK and

recognition, whereas PI-W correlated positively with FOK but was unrelated to recognition.


The contribution of the overall amount of accessible information. As in Experiment 1, FOK judgments increased with the sheer number of letters reported (PI-T). Thus, mean FOK judgments for PI-T values of 0, 1, 2, 3, 4, and 5 were (number of subjects in parentheses) 6.1 ($n = 11$), 13.6 ($n = 16$), 26.2 ($n = 24$), 43.4 ($n = 24$), 58.6 ($n = 22$), and 49.8 ($n = 17$), respectively. In parallel, recognition performance also increased with increasing PI-T, the respective means being 20.8, 32.4, 38.2, 57.1, 73.1, and 89.8, respectively. It is clear, then, that the mere number of letters recalled is a very powerful predictor of recognition memory.

The majority of the letters reported (88.8%) were correct, and accordingly the mean variance of PI-C (1.24) was markedly higher than that of PI-W (0.33). The mean within-subject correlation between PI-T and PI-C was .85, $t(23) = 37.28, p < .0001$, whereas that between PI-T and PI-W was .04, ns . The mean correlation between PI-C and PI-W was $-.46, t(23) = 13.79, p < .0001$. Thus, PI-T serves as a valid predictor of recognition memory, simply because most of its variance is due to PI-C.

Note that subjects in Experiment 2 reported an average of 2.97 letters, that is, 59.4% of the letters presented, compared with 73.9% in Experiment 1. Nevertheless, the mean within-subject correlation between PI-C and PI-W remained high. Apparently, this negative correlation does not stem only from the fact that PI-C sets an upper limit on PI-W but also because factors that increase correct recall reduce commission errors (see the General Discussion section).

Figure 11 presents the mean number of correct and incorrect letters reported for targets that were subsequently recognized and for those that were not. Targets that were ultimately recognized were associated with the report of a larger number of correct letters, $r(23) = 8.88, p < .0001$, and a smaller number of incorrect letters, $t(23) = 6.91, p < .0001$, than targets that were not recognized. Note that in this latter analysis, a negative correlation between PI-W and recognition emerged after all, possibly because of the larger number of observations on which each of the two means was based. Despite the conflicting effects of PI-C and PI-W, however, PI-T was correlated positively with

recognition accuracy, $t(23) = 7.27, p < .0001$. Also, as in Experiment 1, the percentage of correct letters among those reported was larger for recognized (93.6%) than for nonrecognized targets (80.1%),



PI - C




Figure 10. Mean correct recognition plotted (top panel) as a function of correct partial information (PI-C), with wrong partial information (PI-W) as a parameter, and (bottom panel) as a function of PI-W, with PI-C as a parameter (Experiment 2).

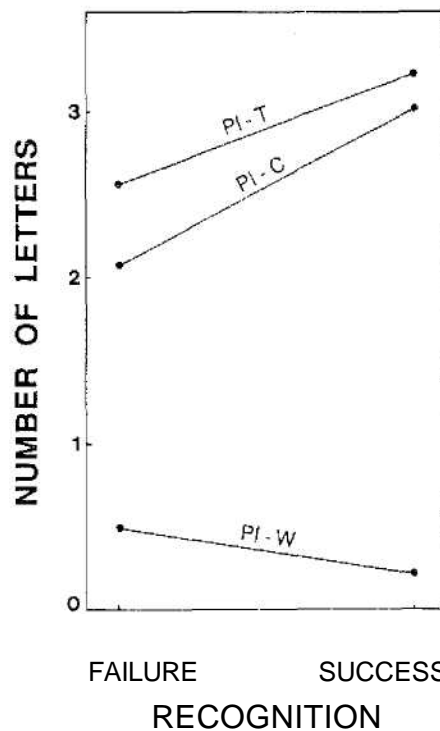


Figure 11. Mean correct partial information (PI-C), wrong partial information (PI-W), and total number of letters reproduced (PI-T) for recognized and not recognized targets (Experiment 2).

Monitoring the accuracy of the accessed information. I now turn to the results pertaining to ease of access, showing how ease of access can be used by FOK to monitor the accuracy of the information retrieved. Three hypotheses were tested: FOK judgments do monitor the accuracy of the accessible information, ease of access is diagnostic of report accuracy, and FOK judgments increase with increasing ease of access.

The left panel of Figure 12 presents mean FOK judgments as a function of PI-T and report accuracy. For the data presented in this figure, all items with the same PI-T value were divided into those where report accuracy was perfect and those where it was less than perfect ($PI-W > 0$). The mean FOK judgments were then calculated for each subject for each of the available cells, and these means were averaged across all subjects for which data were available. It can be seen that FOK judgments generally increase with the amount of partial information (PI-T), as discussed earlier, but that for each value of PI-T, they also increase with the accuracy of the reported letters. A two-way ANOVA (with unequal ns) yielded $F(4, 23) = 22.14, p < .0001$, for PI-T, $F(1, 23) = 13.39, p < .005$, for report accuracy, and $F(4, 4g) = 1.05, ns$, for the interaction.

Ease of access and the accuracy of partial information. I now show that ease of access is diagnostic of the accuracy of partial information. Table 3 presents mean latency to initiate recall as a function of PI-C and PI-W. In calculating these means, all response times below 250 ms or above 15 s were eliminated (1.29%). Response times were then averaged for each subject for

Table 3
Mean Recall Latencies for Each of the FI-C and PI-W Combinations and the Number of Subjects (Experiment 2)

PI-W	PI-C											
	0		1		2		3		4		5	
	M	n	M	n	M	n	M	n	M	n	M	n
0		11	4,994	15	3,201	24	2,232	24	1,925	21	1,598	16
1	8,010	7	4,858	21	2,934	21	3,007	16	2,713	11		
2	5,114	5	3,479	15	2,313	8	1,904	3				
3	3,910	1	2,948	1	3,418	3						
4	—	0										
5	—	0										

jVo/r^l. PI-C = correct partial information; PI-W = wrong partial information.

each of the cells, and these means were then averaged across all subjects.

The results disclose two trends. First, recall latency decreased markedly with the amount of partial information accessible: Mean latencies for PI-Ts of 1, 2, 3, 4, and 5 were (number of subjects in parentheses) 5,622 ($n = 16$), 3,642 ($n = 24$), 2,492 ($n = 24$), 2,5245 ($n = 22$), and 1,798 ms ($n = 17$), respectively, $F(4, 75) = 22.79$, $p < .0001$. Second, ease of access was also associated with the accuracy of partial information: When the total amount of information was held constant, more accurate reports were generally associated with shorter latencies than were less accurate reports. This can be seen in the right panel of Figure 12, which presents ease of access data in a manner that facilitates comparison with the FOK data. The results generally parallel those obtained for FOK judgments: Response latency decreased with the amount of information recalled, but for each PI-T value, it also decreased with the accuracy of the recalled information. A two-way ANOVA (with unequal n s) yielded $F(4, 75) = 20.07$, $p < .0001$, for PI-T, $F(1, 23) = 29.78$, $p < .0001$, for report accuracy, and $F(4, 47) = 2.18$, $p < .10$, for the interaction. Thus, recall latency is a valid cue for the accuracy of the information retrieved (proposition 12).

Ease of access and the FOK. Now I examine the hypothesis that recall latency is actually used by subjects in the computation of FOK. In view of the correlations between PI-T, on the one hand, and recall latency and FOK judgments, on the other, it was imperative to control for PI-T in evaluating the possible dependence of FOK judgments on recall latency. The recall latencies of each subject were split at the median. Then the mean FOK judgments associated with above-median ("slow") and below-median ("fast") latencies were calculated for each value of PI-T. These means, averaged across subjects, are plotted in the left panel of Figure 13. It can be seen that FOK judgments generally increase with increasing PI-T, but that for four out of the five PI-T values they are higher for the fast than for the slow reports.

Each of the points in Figure 13 is based on a different subgroup of subjects for which the respective data were available, and, furthermore, there is a risk of a systematic bias in the kind of subjects contributing to each point. Therefore, the results were also analyzed using a second method. First, for each sub-

ject, mean FOK judgments for slow and fast reports were calculated for each value of PI-T. Then only those PI-T values were retained for which mean FOK judgments were available for both levels of response latency (i.e., above and below the subject's median). These cell means were then averaged for each subject to obtain a mean response latency for slow and fast reports. These means, averaged across all subjects (data were computable for all subjects), yielded 38.82 for the slow reports and




Figure 12. Mean feeling-of-knowing (FOK) judgments (left panel) and mean recall latency (right panel) as a function of recall accuracy for each value of PI-T (total number of letters reproduced; Experiment 2).

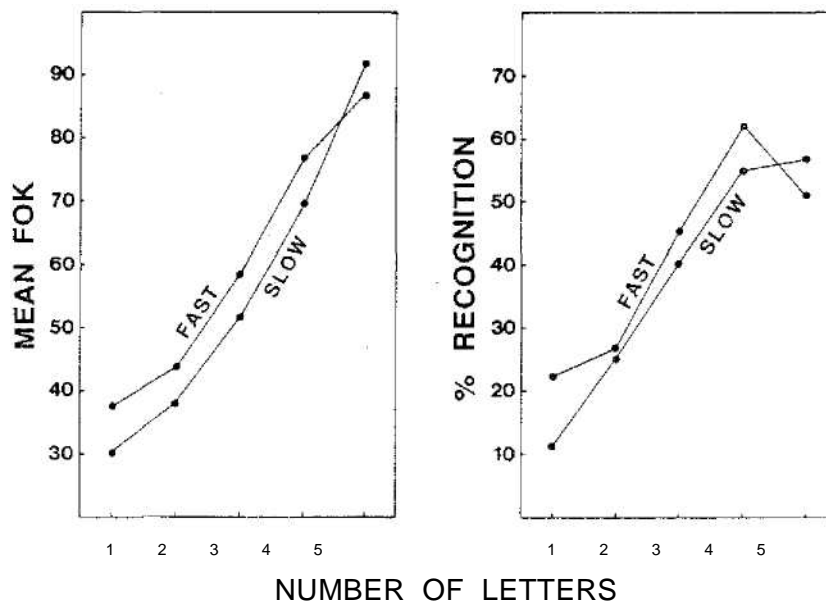


Figure 13. Mean feeling-of-knowing (FOK) judgments (left panel) and mean correct recognition (right panel) for slow and fast reports as a function of the number of letters reported (Experiment 2).

43.10 for the fast reports, $t(23) = 2.76, p < .02$. Thus, FOR judgments increased with decreasing latency of recall even when the number of letters reported was held constant.

Ease of access and recognition performance. The right panel of Figure 13 also presents the mean recognition performance for fast and slow reports as a function of PI-T. It can be seen that recognition memory is somewhat better for fast than for slow partial reports, even when PI-T is controlled. A similar analysis as that reported earlier for the relationship between response latency and FOK judgments yielded 56.63 correct recognition for the slow reports and 61.93 for the fast reports. Although the difference was in the predicted direction, it failed to reach significance, $t(23) = 1.43, p < .10$, one-tailed.

The effects of retention interval on memory and monitoring. The results pertaining to retention interval were only partly consistent with predictions, presumably because the manipulation of memorability was not sufficiently strong. The short retention interval resulted in more correct letters recalled (2.73) than the long interval (2.55), $t(23) = 2.74, p < .02$, but the two intervals did not differ in the number of incorrect letters reported (0.33 and 0.34, respectively), $t(23) = 0.33, ns$. Also, the shorter retention interval was associated with shorter recall latencies ($M = 2,566$ ms) than the longer interval ($M = 2,847$ ms), $t(23) = 2.02, p < .06$, and yielded higher FOK judgments (42.23 and 39.09, respectively), $t(23) = 2.52, p < .02$. However, the manipulation of retention interval did not succeed in affecting recognition memory, the respective means being 57.29 and 57.92,

This latter result, perhaps, explains why the manipulation of retention interval failed to affect monitoring accuracy. Thus, when all trials were classified in terms of FOK judgments (below vs. above the subject's median) and retention interval (long vs. short), a two-way ANOVA, with recognition memory as the

dependent variable, yielded $F(1, 23) = 47.24, p < .0001$, for FOK level, and $F < 1$ for both retention interval and the interaction.

Also, the introduction of differences in retention interval in Experiment 3 would have been expected to result in a higher FOK accuracy than in Experiment 1, because it is expected to increase the interitem variance in memory strength. The results, however, failed to support this expectation.

In conclusion, the results of Experiment 2 generally replicated the findings from Experiment 1. In addition, they provided evidence consistent with the hypothesized role of ease of access in mediating the relationship between FOK and the accuracy of the accessed information. Thus, FOK judgments seem to function in both capacities, as predicting the accuracy of future memory performance (prospective monitoring) and as postdating the accuracy of the information that has already been accessed (retrospective monitoring).

The results on retention interval, however, were only partly consistent with predictions, possibly because the manipulation of memory strength was not sufficiently strong. Retention interval affected the number of correct letters recalled, as well as recall latency and FOK judgments. However, it had no effect on recognition performance and also failed to yield the expected effect on FOK accuracy. This latter failure is considered in the General Discussion section.

Experiment 3

In Experiments 1 and 2, the amount of accessible information was defined in terms of the number of individual letters that could be recalled. However, in many situations, the partial information accessible pertains to features or attributes of the target as a whole, rather than to fragments of the target (see, e.g.,

Durso & Shore, 1991; Yavutz & Bousfield, 1959). The aim of Experiment 3 was to demonstrate that the main predictions of the accessibility account also hold for such situations. Specifically, subjects memorized the Hebrew translations of so-called "Somali" words. When unable to recall the translation, they judged the semantic connotation of the word on the evaluative (good-bad) dimension and provided FOK judgments. Unlike in the previous experiments, in Experiment 3, FOK and attribute judgments were solicited only when the subject failed to supply any response. (Soliciting attribute judgments following commission errors had proved both useless and confusing in exploratory work.) The design of the experiment shares some features with the studies of Schacter and Worling (1985) and Eysenck (1979) but permits measurement of the amount (rather than only the accuracy) of partial information.

Method

Subjects. Twenty-four Hebrew-speaking University of Haifa students participated in the experiment, 2 received course credit, and 22 were paid for their participation.

Stimulus materials- A list of 48 Hebrew words was compiled, 24 having a "good" connotation and 24 having a "bad" connotation. The words were selected on the basis of preliminary ratings, and the final set was submitted to five judges who rated the words on the semantic differential scale of good-bad. Mean ratings were 6.04 and 1.63 for the good and bad words, respectively (with higher scores indicating more positive connotations).

The 48 Somali words were one- to three-syllable pronounceable nonsense strings. They were chosen so that they evoked little definite associations among Hebrew speakers. These were randomly paired with the Hebrew words for each subject, with the restriction that across all subjects each Somali word was equally paired with good and bad words.

Procedure The experiment was conducted on a VAXlab microcomputer. Stimuli were presented on a graphic display unit. The subject's responses were vocal and were recorded by the experimenter from another terminal.

The experiment consisted of a study and a test phase. In the first block of the study phase, the Somali-Hebrew pairs were presented, each for 5 s, with a 1-s interval between pairs. (The Somali words were written in Latin letters.) In the remaining study blocks, only the Somali word was shown, and the subject had to say aloud the Hebrew translation. Eight seconds were allowed. The corresponding Hebrew word was then presented next to the Somali word (for 2.5 s) either 1 s after the subject responded or after 5 s if no response was supplied. There was a 1-s intertrial interval. Seven subjects were presented with three blocks, whereas the remaining subjects required an additional block because they failed to reach the criterion of 33% recall on the third block. Presentation order was random for each subject and block.

There was a 10-min study-test interval during which subjects were given two filler tasks. The first block of the test phase was similar to the study blocks except that only the Somali words were presented. When subjects failed to supply an answer within the 8-s time limit, they were asked to rate the Hebrew meaning of the Somali words as having a good connotation or a bad connotation but to abstain from responding when they had no articulate feeling about the connotation. The words *good*, *abstain*, and *bad* appeared on the screen, and after the subject said aloud one of them, they were replaced by the FOK question, "What are your chances of recognizing the correct target? (17%-100%)." The instructions given at the beginning of the experiment specified the nature of the recognition test and indicated that 17% represents chance performance.

The second test block involved recognition memory. On each trial,

one Somali word appeared on the screen, and beneath it appeared six Hebrew alternatives. The distractors were chosen randomly from the study list, except that across the entire test, each Hebrew word appeared equally often as a distractor and was not repeated on two successive trials. Also, the six alternative answers (including the target) included three good words and three bad words.

Results and Discussion

The number of words recalled averaged 2.83, 8.21, and 11.76 for the three last blocks of the study phase. In the recall test block, subjects recalled the correct word in 19.7% of the trials. In addition, they made a commission error in 8.0 trials, of which 6.7 were intralist intrusions.

The following analyses were confined to the remaining items, that is, those for which subjects made a good-bad-abstain judgment and provided FOK judgments. The number of such items ranged from 7 to 28 across subjects ($M = 20.3$). Subjects provided attribute (good-bad) judgments in 52.5% of the trials (range = 16.7%—81.0% across subjects) and abstained in the remaining trials.

Table 4 presents the mean number of good, bad, and abstain responses as a function of the category of the target word (good or bad). It can be seen that when attribute judgments were made, they were correct in 67.0% of the cases (range = 25%–100%). This performance was significantly better than chance, $t(23) = 4.09$, $p < .0005$, indicating that subjects were able to access the connotative meaning of unrecalled words.

This conclusion is also supported by an analysis of the type of commission errors made in the recall test block. A total of 192 commission errors were made across all subjects and items. Of these, 161 were words from the list and therefore could be readily classified as good or bad. These were distributed as follows: When the correct answer was a good word, the incorrect answer was also a good word in 56.9% of the cases and a bad word in the remaining cases. In contrast, when the correct answer was a bad word, the incorrect answer was a good word in only 38.2% of the cases and a good word in the remaining cases. These results are consistent with those of Yavutz and Bousfield (1959) and Schacter and Worling (1985).

Turning next to FOK judgments, it can be seen (Table 5) that these were higher when attribute identification was correct than when no partial attribute information was provided, (abstain), $t(23) = 5.98$, $p < .0001$. However, FOK judgments were also significantly higher when attribute identification was incorrect, relative to the abstain condition, $t(20) = 2.91$, $p < .001$, and did not differ from those found for correct attribute identifications, $t(20) = 0.67$, *ns*. Note that the latter two comparisons were based on only 21 subjects, because 3 subjects did not make any incorrect good-bad judgments. (The correct and abstain FOK means for these 21 subjects were 48.1 and 31.6, respectively.)

Examination of the means for individual subjects indicated that 22 out of the 24 subjects assigned higher FOK judgments to items for which they provided good-bad judgments than to abstain items ($p < .0001$ by a binomial test). In contrast, of the 21 subjects who made both correct and incorrect identifications, 12 subjects gave higher FOK judgments to items for which attribute identifications were correct than to those for which they were wrong (*ns* by a binomial test).

Table 4
Mean Number of Good, Bad, and Abstain Responses as a Function of the Category of the Target Word (Experiment 3)

Attribute judgment	Target category		
	Good	Bad	All
Bad Good	1.88	2.58	4.46
Abstain	3.83	2.08	5.91
All	4.42	5.45	9.87
	10.13	10.13	20.25

Finally, I also examined the relationships between attribute judgments and FOK judgments, on the one hand, and recognition performance, on the other hand. Somewhat surprisingly, these relationships were very weak. As far as partial information is concerned, percentage of correct recognition was higher for items with correct attribute identification (60.2%) than for items with incorrect attribute identification (53.7%), but the difference was not significant, $t(20) = 1.21$, *ns*. Items eliciting abstain responses resulted in 59.6% correct recognition. Thus, apparently recognition performance was adversely affected by incorrect attribute identification but was not facilitated by correct attribute identification.

FOK judgments were also only weakly predictive of recognition performance: When FOK judgments were split at the subjects' medians, correct recognition averaged 60.6% for items with above-median FOK, compared with 57.3% for items with below-median FOK, $f(23) = 0.87$, *ns*. FOK judgments also failed to predict the selection of recognition foils that were consistent with the connotation (good or bad) of the target: The likelihood of choosing such a foil (including the correct target) was .77 for items with above-median FOK and .75 for items with below-median FOK, $2(23) = 0.41$, *ns*.

In summary, the results indicate that the accessibility account can be extended to situations where the partial information accessible pertains to dimensional attributes of the target, rather than to its fragments. Thus, subjects were more confident about the future recognition of an unrecalable target when they could access attribute information about it, and this was true irrespective of whether that information was correct or not. However, FOK judgments in the present experiment were not predictive of recognition memory (see also, e.g., Schwartz & Metcalfe, 1992), preventing evaluation of some of the model's assumptions regarding the basis of FOK accuracy.

General Discussion

In this article, I explored three questions about the FOK within a unified model: First, what is the validity of FOK in predicting memory performance, second, what is the process underlying FOK judgments, and, third, how does this process account for the accuracy of FOK predictions? Although an attempt was made to address each of the questions separately, it is clear from the model proposed, as well as from the empirical data, that the three questions are intimately linked and ought to

be considered conjointly within a common conceptual framework.

The following discussion is divided into five sections. The first concerns the basis for the FOK, whereas the second focuses on explaining the predictive validity of FOK. The third section examines the accessibility account in relation to other accounts, and the fourth section takes up the general issue of privileged access to an internal monitor. Finally, the factors contributing to FOK accuracy and inaccuracy are discussed.

The Basis of the FOK

The view of FOK advanced in this article stands in contrast with the trace-monitoring metaphor, which has implicitly guided much of the prevalent theoretical formulations and methodological practices in the study of FOK. This metaphor assumes independence between a monitoring process that directly detects the presence of the memory target in store and a retrieval process that attempts to recollect it. The view of FOK as "a storage state indicator" (Hart, 1967a, p. 689) explains, perhaps, why most previous research has concentrated on demonstrating the accuracy of FOK rather than on explaining it, because, in that view, the validity of FOK stems directly from the trace-access postulate. In fact, the implicit assumption is that FOK's validity constitutes, in itself, evidence that subjects can detect the availability in store of an otherwise inaccessible target (to use the terminology of Tulving & Pearlstone, 1966).

In contrast with the view of FOK as monitoring memory storage, the position I advance in this article assumes that FOK judgments primarily monitor the retrievability of information pertinent to the solicited target. This view is similar to what Nelson and Narens (1990) referred to as the No-Magic Hypothesis: "The FOK does not directly monitor a given unrecalled item in memory, but rather the FOK monitors recallable aspects related to that item, such as the item's acquisition history or partial/related recalled components" (p. 158). They added, "Notice that this way of making FOK judgments would utilize only suprathreshold information about remembered attributes of the item (including incorrectly remembered suprathreshold information!)" (p. 158). Thus, according to the accessibility position, FOK is based on the outcome of the very process used to search for the target and retrieve it from memory. A memory pointer that leaves one absolutely "blank" produces a feeling of not knowing (see Gruneberg, 1978; Kolers & Paley, 1976). In

Table 5
Mean Feeling-of-Knowing Judgments and the Number of Subjects for Correct Attribute Identification, Incorrect Attribute Identification, and Abstain Responses (Experiment 3)

Variable	Attribute judgment		
	Correct identification	incorrect identification	Abstain
<i>M</i>	50.1	47.3	34.9
<i>n</i>	24	21	24

contrast, when a pointer activates some clues, these clues can be used in computing FOK when recall of the target fails. This computation sometimes entails a conscious, deliberate calculation of probabilities (see the following paragraphs), but in other cases it is the sheer accessibility of information that counts: If the aborted process leaves behind well-articulated partial information, such as fragments of the target, or beliefs about prior encounters (see Costermans et al., 1992; Nelson et al., 1984; Nelson & Narens, 1990), it is likely to give rise to a positive FOK and even to the subjective experience that recall is imminent (see R. Brown & McNeill, 1966).

The important feature of the accessibility account lies in the assumption that FOK judgments have no access to information about the solicited target that is not already contained in the outcome of the retrieval attempt. If the search process goes astray, so will the FOK. Clearly, in some cases, the information stored in the subject's memory may simply be wrong (e.g., the subject believes that Sydney is the capital of Australia), so that the subject's target differs from the experimenter's intended target (see R. Brown & McNeill, 1966; Koriat & Lieblich, 1974; Nelson et al., 1984). In such cases, it would appear odd to assume that FOK judgments monitor the presence of the correct target in store when the subject's search is guided by a different referent (see Nelson & Narens, 1990). However, even when the person does possess the correct information, the search process may be misled by a variety of spurious activations, such as those emanating from neighboring targets (see, e.g., A. S. Brown, 1991; Jones & Langford, 1987; Koriat & Lieblich, 1977; Mozer, 1991). Furthermore, metacognitive judgments are affected not only by information that is specific to the target in question but also by knowledge of the characteristics common to the class of items to which the target belongs (e.g., Glenberg, Sanocki, Epstein, & Morris, 1987; Koriat & Lieblich, 1974, 1977). For example, one may feel that one knows the capital of Sweden because one can recall the capital of Denmark (Gruneberg, 1978). Presumably, it is particularly in the early stages of the retrieval process, when the search is still unfocused, that FOK judgments tend to be based on class information (see Koriat & Lieblich, 1974) or activation from remote associates of the target.

Many of the previous FOK studies have been implicitly guided by a trace-access view of FOK, and their methodology does not permit a fair assessment of the accessibility account. As mentioned earlier, this methodology implicitly assumes that subjects know that they know. Thus, first, the practice to solicit FOK judgments only when subjects fail to retrieve the correct target assumes that when subjects get the correct answer, they also know that it is correct. Second, the methodology used in previous work on the partial information hypothesis, which focuses on the accuracy of the information retrieved, also implies that subjects actually know how much they know.

In the present study, both of the biases mentioned earlier were eliminated from the design of Experiments 1 and 2 by soliciting FOK judgments regardless of recall performance and by using a free-report testing procedure. In Experiment 3, FOK judgments were solicited only following recall omission, but when omission occurred, subjects were still given the option to make attribute judgments or refrain from doing so.

The results obtained with these procedures were clearly in support of the accessibility position. In both Experiments 1 and 2, FOK judgments were positively correlated with the amount of accessible information as such, whether that information was correct or not. Thus, the critical factor affecting FOK was the sheer accessibility of information, not its accuracy. A similar result was also found in Experiment 3. Furthermore, the links between FOK judgments, partial information, and recognition in Experiments 1 and 2 also supported the idea that PI-C contributes to the success of FOK in predicting recognition performance, whereas PI-W generally contributes to its failure.

One previous finding that accords well with the present results and with the accessibility model is that FOK judgments are significantly higher following commission errors than following omission errors (e.g., Krinsky & Nelson, 1985; Nelson et al., 1984). Clearly, a commission error is evidence that some information could be retrieved about the target, and therefore FOK judgments ought to be higher here than following omission.

The accessibility account implies a continuum of retrieval states from those where a pointer brings nothing to mind to those that culminate in complete recall, correct or incorrect. It may be asked whether complete recall might not represent a qualitatively different state, boosting FOK judgments beyond what would be expected from the overall relationship between FOK and the amount of information retrieved. The results do not support this possibility: In Experiment 1, mean FOK judgments for PI-Ts of 0, 1, 2, 3, and 4 were 17.9, 20.8, 40.7, 63.7, and 87.6, respectively. The means for PLTs 0-5 in Experiment 2 were 6.1, 13.6, 26.2, 43.4, 58.6, and 49.8, respectively.

Explaining the Accuracy of FOK Judgments

Let me turn next to examine the basis for the accuracy of FOK judgments. From the trace-monitoring view, one would expect FOK judgments to be quite accurate in predicting recognition memory (see Costermans et al., 1992; Reder & Ritter, 1992). Thus, curiously enough, a preoccupation, with explaining the inaccuracy of FOK judgments (see Nelson et al., 1984; Nelson, Leonesio, Landwehr, & Narens, 1986) is often symptomatic of the implicit endorsement of a trace-monitoring view of FOK. In contrast, if the accessibility view is accepted, then the primary challenge becomes that of explaining the accuracy of FOK judgments.

The evidence reported in this article supported both of the following propositions: first, that FOK is indeed based on the mere accessibility of information and, second, that it is nevertheless generally predictive of subsequent memory performance. In explaining this pattern, it was proposed that the efficacy of metamemory processes stems from the quality of memory processes themselves. Thus, the answer to the question of why FOK judgments are accurate ought to be found in the effectiveness of memory storage and retrieval, not in the accuracy of some specialized structure that is dedicated to the monitoring of memory storage (e.g., Hart's MEMO). This assumption, of course, does not deny either the reality or the importance of memory-monitoring processes. It implies, however, that the accuracy of FOK should vary greatly with the many

factors that affect retrieval (see Leonesio & Nelson, 1990; Lupker et al. 1991).

Why, then, are FOK judgments generally accurate? The results of the present study suggest three reasons. The first, and most important, lies in the overall accuracy of the accessible information. In general, items committed to memory are more likely to give rise to correct than to incorrect recall. As noted earlier, what matters is not the likelihood that a memorized item is accessible (input-bound performance) but the likelihood that an accessed item is correct (output-bound performance; see Koriat & Goldsmith, 1993a). This is because FOK judgments are assumed to be based on the output of the retrieval attempt. Indeed, in the present study, a reported letter had a .89 probability to be correct in both Experiments 1 and 2. In Experiment 3, when subjects provided a (complete) answer, that answer was correct in .71 of the cases (counting as wrong both intralist and extralist commission errors). When no answer was provided and an attribute judgment was volunteered, that judgment was correct in .67 of the cases. Thus, on the whole, an accessible answer stands a better chance of being right than wrong. Therefore, PI-C constitutes the largest bulk of PI-T and makes a greater contribution to its variance. (The variance of PI-C was about three to four times larger than that of PI-W in Experiments 1 and 2.)

Furthermore, examination of the intercorrelations between PI-C, PI-W and PI-T indicates that the variance of PI-T was determined practically entirely by the variance of PI-C: In Experiment 1, mean within-subject correlation between PI-T and PI-C was .83, whereas that between PI-T and PI-W was .00. The respective correlations for Experiment 2 were .85 and .04. This striking pattern probably stems from the fact that the relationship between PI-W and PI-T was suppressed by the negative correlation between PI-W and PI-C. Thus, although PI-W is possibly responsible for the inaccuracy of FOK judgments, its independent contribution to PI-T is very small.

The second reason is that the amount of information available (PI-T) is also diagnostic of its accuracy. The data on which this conclusion is based have not been reported so far and are presented now. In both Experiments 1 and 2, there is a trend suggesting that the larger the number of letters reported, the better the likelihood that each of them is correct. Report accuracy, defined as the proportion of correct letters among those reported (see Koriat & Goldsmith, 1993a), averaged .85, .83, .86, and .92 for reports of one, two, three, and four letters, respectively, in Experiment 1. The respective means for Experiment 3 were somewhat more orderly: .74, .83, .89, .90, and .93 for reports of one to five letters. One-way ANOVAs on these means yielded $F(3, 74) = 3.96, p < .05$, for Experiment 1 and $F(4, 75) = 5.33, p < .005$, for Experiment 2. Thus, the number of letters recalled is not only diagnostic of the strength of the memory trace of the target as a whole but also of the accuracy of each piece of partial information retrieved.

The third reason has to do with the quality of retrieval itself, that is, the intensity of the accessible information. The recall latency data of Experiment 2 indicated that ease of access is diagnostic of the accuracy of the information retrieved and that FOK judgments increase with increasing ease of access. There was also a trend suggesting improved recognition memory for

items associated with shorter recall latencies than for those associated with longer latencies. Thus, reliance on ease of access can contribute to FOK validity in predicting recognition.

These results have some bearing on retrospective monitoring as well. Thus, the results pertaining to the quantity-accuracy correlation, and those concerning ease of access, suggest two vehicles by which subjects can monitor the accuracy of the information that has been retrieved. These results can explain why Schacter and Worling's (1985) subjects were accurate in monitoring their attribute judgments about unrecalable words (although this was not true in Experiment 3 of the present study) and why FOK judgments in both Experiments 1 and 2 were affected by the accuracy of the partial information recalled. Indeed, the results obtained by Nelson and his associates (Nelson et al., 1984; Nelson, Dunlosky, White, Steinberg, Townes, & Anderson, 1990) indicate that subjects can monitor the correctness of their answers and that this success may be mediated, in part, by the longer recall latencies for correct recalls than for commission errors. Similar results have also been reported more recently by Costermans et al. (1992).

In summary, the results on the whole suggest that it is not necessary to invoke the notion of trace access to account for the validity of FOK. Rather, subjects can effectively monitor the correctness of their past and future responses on the basis of the mere accessibility of information. Thus, the amount and intensity of the information retrieved can serve as potent cues for prospective and retrospective judgments regarding the accuracy of their memory performance. These cues would be generally valid as long as the partial information retrieved is more likely to be correct than false. When such is not the case (e.g., Fischhoff et al., 1977; Koriat, 1976; Koriat & Lieblich, 1977), FOK judgments would not be expected to be predictive of successful memory performance. Hence, the accuracy of metamemory is intimately linked to the accuracy of memory itself.

The postulated link between memory and metamemory implies that manipulations that improve memory performance should also improve metamemory accuracy (proposition 11 of the model). However, previous studies that examined this relationship have yielded inconsistent results (Carroll & Nelson, 1993; Carroll & Simington, 1986; Lupker et al., 1991; Nelson et al., 1982; Nelson & Narens, 1990; Schacter, 1983; Schacter & Worling, 1985; Schwartz & Metcalfe, 1992). The results of the present study were also ambiguous in that, on the one hand, the manipulation of retention interval in Experiment 2 failed to affect FOK accuracy, but, on the other hand, the analysis of individual differences in Experiment 1 (see Figure 7) indicated that FOK accuracy was better for high-recognition than for low-recognition subjects.

In view of these discrepancies, it is important to specify the model's predictions regarding the link between memory and metamemory. Although manipulations that improve memory are expected to increase the means of both FOK and recognition scores, such manipulations may sometimes reduce their variances, thereby impairing monitoring accuracy. This is because monitoring accuracy is commonly defined in terms of the within-subject correlation between FOK and recognition (see also Nelson et al., 1986; Nelson & Narens, 1990). Therefore, if all of the items in a test are uniformly "easy," subjects may find

it difficult to predict the small differences in the likelihood of correct recognition on the basis of the small differences in accessibility. In the extreme case, when a subject demonstrates perfect recognition performance, and makes positive FOKs on all unrecalled answers (thereby being, in some sense, perfectly accurate), the FOK-recognition correlation will simply be non-computable, because of the uniformity of the scores. Because this correlation is very sensitive to the variance in the overall "difficulty" of the items that enter into the computation, it should be expected to be higher when FOK judgments are made on all items than when they are made on only unrecalled items (see Reder & Ritter, 1992).

In general, according to the accessibility model an experimental manipulation that improves memory performance should improve metamemory accuracy only when it also increases PI-C, or more precisely, when it increases the proportion of P.I-T variance that is due to PI-C. This condition is satisfied in the comparison between the performance of low-recognition and high-recognition subjects (see Figure 7), which demonstrated the expected memory-metamemory link.

The Accessibility Account of FOK in Relation to Other Accounts

Although in the present article I focus on the accessibility of partial information as the basis for the FOK, other cues may also enter into the computation of FOK, consistent with the view of FOK as being multiply determined (see Lconcsio & Nelson, 1990). Nelson et al. (1984), for example, presented a comprehensive review, listing 12 possible mechanisms, broadly classified into two groups: trace-access mechanisms and inferential mechanisms. In this section, I attempt to clarify the relation between accessibility and three mechanisms discussed in the literature: target retrievability, inference-based mechanisms, and cue familiarity.

Consider target retrievability first. This mechanism, designated "partial recall" by Nelson et al. (1984), has been classified by them as a "trace-access mechanism" because FOK is assumed to rest on gaining partial access to the target proper. As noted earlier, this type of mechanism is also implied by those researchers who focused on the accuracy of partial information as the basis for FOK (Blake, 1973; Eysenck, 1979; Schacter & Worling, 1985). Furthermore, it is this mechanism that is generally intended in recent discussions that contrast the effects of cue priming on FOK with those of target priming (see Metcalfe, 1993; Schwartz & Metcalfe, 1992).

Thus, in both the accessibility mechanism and the target-retrievability mechanism, FOK is seen to monitor retrieval rather than storage. However, unlike the former mechanism, where FOK is seen to rest on the mere accessibility of partial information, in the latter mechanism FOK is assumed to be specifically tuned to the partial recall of the actual, correct target. This implies that subjects can monitor directly the accuracy of the information that comes to mind.

Second, it is important to clarify the relationship between the accessibility mechanism and the class of inferential mechanisms discussed by Nelson et al. (1984). As noted earlier, subjects can monitor indirectly the dependability of the retrieved

information, and they do so on the basis of two types of cues, those that pertain to the intensity of the accessible information and those that have to do with the content of the information. Although I focus on only the former cues in this article, it is important to briefly examine the contribution of content considerations (see also Nelson et al., 1984).

It is proposed that in the early stages of the search process, FOK is based on a shallow analysis of the memory pointer and is affected by the amount and intensity of activations regardless of their source. In many situations, however, subjects also consult the content of the information retrieved. Such is often the case with real-world knowledge, particularly during the later stages of the search process. For example, one may judge the retrieved information to be implausible on the basis of previously retrieved clues. Alternatively, one may identify the source of a temporary activation and discount its influence by attributing it to that source (see Jacoby & Kelley, 1987; Jacoby & Whitehouse, 1989). In such situations, a retrieved piece of information may, in fact, reduce rather than enhance FOK judgments.

When the content of the retrieved information enters into consideration, the monitoring process changes its quality from an automatic, nonanalytic process to a deliberate, inferential process of probability estimation (see Jacoby & Brooks, 1984; Jacoby & Kelley, 1987). Possibly many of the FOK judgments provided by subjects in FOK studies actually represent educated guesses that are based on an analytic process rather than on an immediate feeling. In fact, it would be difficult to think of some questions as capable of eliciting a FOK about the unrecalled answer (e.g., "How old was Dwight Eisenhower when he died?"), though subjects can still make predictions regarding their chances of selecting the answer among distractors. In such cases, subjects may prefer to phrase their judgment as "I believe that I know the answer" or "I ought to know the answer" rather than as "I feel I know it" (see also Costermans et al., 1992). One may question the inclusion of such clearly calculated inferences under the rubric of FOK,

Of course, whether or not content considerations enter into the computation of FOK depends critically on the nature of the memory pointer. Such considerations possibly play a more restricted role when the memory task is about new, experimentally presented information than when it taps real-world knowledge. This is one reason why the tasks used in the present study are of the former type: They help avoid the complications introduced by considerations having to do with the content of the information accessed. Thus, in Experiments 1 and 2, for example, subjects apparently had little independent clues by which they could judge whether a letter that came to mind was indeed correct or not. Note that this situation is also characteristic of some of the memory pointers that tap long-term, semantic or episodic knowledge. For example, in the TOT state, the information that comes to mind (e.g., fragments of a name or a word) is of a sort that is difficult to pit against other pieces of information. In fact, the TOT state is commonly precipitated with regard to proper names, and these usually have very few semantic connections (see Burke et al., 1991; Cohen, 1990) so that people may find it difficult to judge the plausibility of the fragmentary information that comes to mind.

In summary, in terms of the classification proposed by Nelson et al. (1984), the accessibility heuristic can be clearly distinguished from two types of direct-access mechanisms: The first is where FOK is assumed to directly monitor memory storage (e.g., Hart, 1967a) and the second is where FOK is seen to narrowly monitor the retrievability of the target proper (see Schwartz & Metcalfe, 1992). However, although the accessibility heuristic may be properly classified as inferential, it is important to distinguish between the type of analytic inferences that are deliberately drawn on the basis of content considerations and those that are based strictly on the sheer accessibility and intensity of information.

Finally, I would like to examine the recent work on the cue-familiarity hypothesis, which, on the face of it, appears to be at odds with the present proposal. According to that hypothesis (see Metcalfe, 1993; Nelson et al., 1984; Reder & Ritter, 1992; Schwartz & Metcalfe, 1992), FOK is strictly based on the familiarity of the cue or of the question, not on the retrievability of the answer. This view has been supported by several findings indicating that FOK judgments can be enhanced by advance priming of the cues but not by the priming of the target (Reder, 1987, 1988; Reder & Ritter, 1992; Schwartz & Metcalfe, 1992).

Consider first the effects of cue familiarity on FOK. One possibility is that these effects are actually due to accessibility. Thus, on the one hand, an unfamiliar memory pointer is one that leaves the person "blank," whereas a familiar pointer is one that brings some associations to mind. Reder (1987), for example, suggested that the effects of cue familiarity may be conceptualized in terms of the activations emanating from the terms in the question (see also Reder & Ritter, 1992). On the other hand, perhaps it is the ease with which information comes to mind that serves as the cue for the very experience of familiarity (e.g., Jacoby & Dallas, 1981; Jacoby & Kelley, 1987).

Alternatively, cue familiarity may make an independent contribution to FOK. In that case, perhaps cue familiarity enhances FOK very early in the search process, to the extent of motivating search and retrieval, but thereafter it is the accessibility of information, and in some cases its content, that determines FOK judgments. This proposal would accord well with Mandler's dual-process theory (see Mandler, 1991), which distinguishes between a fast, perceptual process that is based on the sheer sense of familiarity or activation and a conceptual process requiring retrieval and elaboration. It is also consistent with Jacoby's distinction between familiarity and recollection. In fact, the recent work of Jacoby and his associates (see Jacoby et al., 1992) on the effects of priming on perception and memory implies that advance cue priming (and hence increased cue familiarity) should enhance FOK judgments only when the subject fails to attribute the subjective experience of familiarity to its proper source. Thus, cue priming, as such, need not automatically enhance FOK.

It should be noted that the cue-familiarity hypothesis is somewhat difficult to apply to the results of the present study. This is because, unlike the tasks used in studies of the cue-familiarity hypothesis where each target was associated with its own distinctive cue (a question or a word), no cues were explicitly used in Experiments 1 and 2 of the present study (subjects simply recalled the most recent target). Thus, an explanation of the

results of these latter experiments in terms of the cue-familiarity hypothesis would probably require additional assumptions.

Turning next to the finding that target priming does not generally enhance FOK, although these seem to argue against the target-retrievability hypothesis, their implications for the accessibility hypothesis are presently not clear. According to this latter hypothesis, for target priming to enhance FOK, it must result in the cue activating a large amount of information early in the search process, that is, before the complete target is retrieved. (Clearly, if priming is effective enough to enhance immediate recall, and if FOK judgments are solicited regardless of subjects' recall performance, then priming would most likely enhance FOK.) This is likely to occur if target priming also enhances the cue-target association, so that the cue alone is capable of evoking more activations than when the target is not primed. Indeed, such appears to be the case in some of the studies where increased memorability of the target did enhance FOK judgments (e.g., A. S. Brown & Bradley, 1985; Lupker et al., 1991; Nelson et al., 1982; Schwartz & Metcalfe, 1992).

In summary, when the results on the effects of cue and target priming are interpreted in terms of the accessibility hypothesis, they would seem to suggest an interesting idea: If one wishes to enhance the amount of information that comes to mind when a question is initially presented, one is better off manipulating aspects of the question itself than manipulating aspects of the sought-after target.

The Issue of Privileged Access to an Internal Monitor

I turn now to the core issue that was raised in the introductory paragraphs: Can subjects directly monitor the contents of their memory? Do they have privileged access to some internal directory where information about the inaccessible target is directly available? This question is similar to that posed by Nisbett and his associates (Nisbett & Bellows, 1977; Nisbett & Wilson, 1977) regarding the basis for subjective reports. They argued that people have little or no direct introspective access to mental processes such as those affecting judgment and decision. The question is whether the same is true regarding FOK judgments. In a way, the subjective experience associated with a strong positive FOK is often similar to what is implied by the trace-access view: Subjects simply report that they sense the unrecalled target and feel its emergence into consciousness (James, 1893). The question is whether such subjective introspections indeed faithfully mirror the internal processes.

According to the view proposed here, subjects have no privileged access to the presence of information in store but must infer their FOK from the products of their memory. If so, an external observer should also be able to predict the subject's recognition performance on the basis of such memory products. Thus, in line with the methodology of Nisbett and Bellows (1977), investigators can pit the subject's predictions against their own predictions by comparing the predictive validity of the subject's FOK judgments with predictions that are based strictly on the sheer number of letters retrieved. This comparison can reveal whether the subject's FOK judgments have access to some privileged information that is not already contained in the output of the search-and-retrieval attempt.

In the analyses reported in the following paragraphs, recognition performance was predicted from (a) subject's FOK judgments and (b) the number of letters reported. Different methods were used to ascertain that the results are not specific to a particular method of analysis. Consider first the results of Experiment 1: In the first analysis, each subject's FOK judgments were divided at the subject's median, assuming that above-median judgments signify positive FOK and below-median judgments signify negative FOK. The hit rate in predicting recognition across all subjects and items was 61%.

Compare that with a simple algorithm that is based on the number of letters reported. Here negative is defined as two letters or fewer and positive as three or more. This algorithm yielded a 72% hit rate. Thus, there is enough information in the mere amount of information accessible to permit successful prediction of recognition, and this prediction does not fall short of that afforded by the subject's FOK.

The same two-way classification of the responses allows calculation of within-subject gamma correlations. Gamma correlations between recognition and the dichotomized FOK scores averaged .55 ($N = 30$). The respective correlation between the dichotomized number of letters reported and recognition was .58 ($N = 27$). The correlation between dichotomized number of letters and dichotomized FOK scores was .94 ($N = 27$), suggesting that FOK judgments rested practically entirely on the amount of information accessible. (Note that the last two correlations were based on 27 subjects only, because 3 subjects never produced fewer than three letters. When the FOK-recognition correlation was recalculated, eliminating these subjects, it was .51.) Thus, FOK is no better a predictor of recognition than the mere number of letters produced.

A second analysis was based on the mean within-subject point-biserial correlations across the entire range of the number of letters and the entire range of FOK judgments. Both the number of letters and FOK judgments predicted recognition to about the same extent: The average correlations across subject were .33 and .34, respectively. The average correlation between FOK and the number of letters reported was .78.

Similar analyses were carried out on the data of Experiment 2. In the first analysis, subjects' FOK judgments were divided at the median, as earlier, and the number of letters reported were also split between two letters or fewer and three letters or more. The hit rate in predicting recognition across all subjects and items was 59% when FOK was the predictor and 66% when the number of letters was the predictor. Within-subject gamma correlations between the dichotomized FOK judgments and recognition averaged .47, whereas those between the dichotomized number of letters reported and recognition averaged .56. The correlation between dichotomized number of letters and dichotomized FOK was .86. Also, the point-biserial correlations between recognition and the full range of FOK judgments averaged .21 across subjects. The respective correlations between recognition and the full range of number of letters reported averaged .31. Note that the correlation between FOK and the number of letters recalled was .59.

In conclusion, the results presented earlier clearly indicate that the mere number of letters recalled is as good a predictor of future recognition performance as the subject's FOK judg-

ments. Thus, the feeling that one "knows" the target does not seem to be any more diagnostic of the "availability" of the solicited information than the mere amount of information accessed. That is, subjects' FOK judgments do not have access to information that is not already contained in the output of the retrieval attempt.

It should be pointed out that Nelson and Narens (1990) also addressed the issue of privileged access, but their approach was very different from that used here. In their study, while the subject was going through a recall-and-FOK procedure, another subject (observer) observed the subject's performance (his or her face, time spent on the question, and so on) and then made predictions regarding how likely the subject would be to recognize the correct answer. The observer's predictions had above-chance accuracy, but the subject's own FOK judgments were still better predictors, suggesting that subjects have idiosyncratic information that they can use in making FOK judgments (see also Nelson et al., 1986). The present study, in fact, can be seen to reveal just what that idiosyncratic information is: the amount and intensity of the clues that come to mind.

Contributions to the Accuracy and Inaccuracy of FOK

In this final section, I summarize the implications of the accessibility model regarding the factors that affect FOK accuracy. Clearly, the accessibility position predicts a greater variability in FOK accuracy across different conditions than would follow from the trace-monitoring position.

According to the proposed account of FOK, the accuracy of FOK depends on the extent to which the accessibility of information at time 1 is predictive of the accuracy of memory performance (e.g., recognition) at time 2. It follows that the failure of FOK to predict recognition performance can result from two sources: the difference in the *memory property* tapped, accessibility versus accuracy, and the effects that are due to *time lag*.

Before examining the contributions of these factors, I wish to consider a third factor, which is methodological in nature. There is a tendency among FOK researchers to treat the FOK-recognition correlation obtained in their studies as if it was an estimate of some immutable underlying property of memory. However, this correlation depends markedly on the statistical properties of the memory test used. In particular, it depends on the variance among the items in properties that are generally related to item difficulty (see Nelson & Narens, 1990). To return to the illustration used earlier, if the items sampled are uniformly easy, the FOK-recognition correlation will tend to be very low. All it takes to boost this correlation, however, is to compile a second sample of very difficult items (e.g., "what was the maiden name of Columbus's grandmother?"), so that the subject would be likely to make a negative FOK on each item and fail to identify the correct target. If this sample of difficult items is added to the sample of easier items, then the FOK-recognition correlation, calculated across all items combined, will increase substantially. Thus, although "FOK accuracy" is expected to be very low (or, in the extreme case, noncomputable) when estimated on the basis of either sample alone, it may become close to perfect when estimated across both samples.

A second property that is important from the point of view of

the accessibility model is *mean output-bound memory performance* (see Koriat & Goldsmith, 1993a), that is, the distribution of correct and incorrect answers among those that a memory pointer tends to bring to mind. In particular, FOK accuracy requires that the largest portion of variance among the items should be the likelihood of bringing to mind correct partial information.

Let me turn next to the effects of the factors mentioned earlier. Consider first the effects of memory property. The present study has clearly shown that accessibility at time t is a valid predictor of accuracy both at time t and at time I . However, the size of the accessibility-accuracy correlation should depend on several factors. First, apparently partial activations ought to be "suprathreshold," that is, subjectively accessible, to affect FOK. Therefore, many factors that influence memory performance, particularly those associated with implicit memory, may sometimes have little effect on FOK (see Jameson, Narens, Goldfarb, & Nelson, 1990). In fact, Nelson and Narens (1990) presented data indicating that subjects' beliefs about the number of times a target had been recalled influenced FOK more strongly than the actual frequency of previous recalls. Second, if FOK monitors accessibility, then it should prove more effective in predicting recall than recognition. Indeed, Jameson et al. (1990) and Reder and Ritter (1992), who validated the accuracy of FOK, against recall performance, found relatively high correlations.

The final factor concerns the source of memory errors in the particular task used. A detailed discussion of the effects of this factor (e.g., PI-W) is beyond the scope of the present article (see Senders & Moray, 1991), but it is clear that the accessibility model cannot be specified in full without a theory of the etiology of memory errors. It can be shown that the accuracy of FOK depends on how the occurrence of false partial recall is correlated with that of correct partial recall. For example, if errors are more likely to intrude when the target's trace is weak, then the correlations designated by links *b*, *c*, and *g* in Figure 1 would all be expected to be negative, and, under some conditions, PI-T will be practically as good a predictor of recognition as PI-C. In contrast, if for some reason the likelihood of a memory error increases with memory strength, then the accuracy of FOK will be primarily determined by the size of the PI-C variance relative to that of PI-W so that when the variance of PI-W is relatively large, FOK judgments may correlate negatively with recognition.

I now report some illustrative results pertaining to the source of incorrect recalls in Experiments 1 and 2. Many of the commission errors in these experiments appear to represent perseveration errors where the letter or letters reported are carried over from the preceding trial. Thus, in Experiment 1, the probability that an incorrect letter originated from the preceding target was .57, when the expected probability was only .25. The respective values in Experiment 2 were .62 and .33. Importantly, when a perseveration error occurred with regard to a particular item, recognition performance for that item was lower in comparison with items where the incorrect error came from a different source: Thus, for all items where exactly one incorrect letter was reported, recognition performance averaged 40% when that letter appeared, on the preceding stimulus, compared

with 60% when it was not. The respective means for items where two incorrect letters were reported were 23% and 39%, respectively. A similar pattern was also observed in Experiment 2. These results can be interpreted to indicate that perseveration errors are more likely to occur for items with weaker memory traces. If this interpretation is correct, then conditions where perseveration errors are dominant should be associated with a relatively high FOK-recognition correlation.

In general, then, the accuracy of FOK in a particular task depends on the source of memory errors in that task and the conditions for their occurrence.

With regard to the effects of time lag, one source of inaccurate FOKs is the systematic changes that occur over time in both the amount and kind of information accessed. First, FOK accuracy may be expected to decrease with increasing time interval between t and t . In general, FOK judgments should be best predictive of immediate recall, as seems to be the case with the TOT state (R. Brown & McNeill, 1966). Second, although accessibility of information at time t is generally predictive of the accessibility of information at time t , this is not always the case. Notably, there are tasks such as solving insight problems (see Metcalfe, 1986; Metcalfe & Wiebe, 1987), where information does not accumulate gradually but rather the answer appears to pop up suddenly. Here the information accessible before the solution may be of little diagnostic value regarding the prospects for a solution (see Metcalfe, 1986). Similarly, there are memory questions where the target appears to be retrieved suddenly as a result of the spontaneous restructuring or paraphrasing of the question (see Koriat & Liebllich, 1977). Such questions would generally lead to an unwarranted feeling of not knowing.

Third, the information that comes to mind when memory is first queried is likely to differ in quality from that which is ultimately used to support target retrieval. For example, it has been argued that "preliminary" FOK judgments are based on a rapid, shallow analysis of the memory pointer rather than on the detailed evaluation of the information retrieved (see, e.g., Koriat & Liebllich, 1977; Reder and Ritter, 1992). Also, in discussions of the TOT state, it has been proposed that the search of the target involves a gradual narrowing of focus so that broad characteristics of the target are activated before more specific information about the target is accessed (see A. S. Brown, 1991). Thus, blockers or interlopers originating from compelling words at the vicinity of the target must first be suppressed before the target itself can be retrieved. Such systematic changes over time in the type of information accessible during the course of a retrieval episode should generally reduce the overall predictive validity of FOK judgments.

One implication of this idea is that FOK accuracy may be improved by delaying FOK judgments. Indeed, with regard to judgments of learning, Nelson and Dunlosky (1991) observed that when these judgments are delayed until shortly after study, they are more accurate than when they are made immediately after study.

A fourth, related factor concerns the possibility of systematic differences between different memory pointers in the rate of information accrual. This possibility was suggested by Koriat and Liebllich's (1977) analysis of memory pointers (see also Morris,

1990). Typically, pointers that resulted in a discrepancy between knowing and FOK were of two types: those that activate rich associations early in the search process, which later prove ineffective in supporting retrieval, and those that bring to mind few associations initially that are followed later by a spontaneous retrieval of the answer. In view of the reliable differences found between memory pointers in that study, it would be of interest to trace the time course of information accumulation that characterizes different pointers and examine the accuracy of FOK judgments solicited for these pointers at different stages of the search process.

Finally, it is important to distinguish between memory tasks that allow the content of the accessed information to be used in computing FOK judgments and those that do not. As noted earlier, the manner in which the rate of Information accrual affects FOK judgments ought to differ in the two cases, and these differences should also influence FOK accuracy.

In addition to the factors mentioned earlier, there is one serious complication that must be considered when drawing conclusions regarding the accuracy of FOK judgments. This complication derives from the fact that the processes that occur at time t_1 are not merely diagnostic of those that operate at time t_2 but may actually influence such processes. So far I have avoided discussion of the possible causal effects of FOK on cognitive processing (proposition 3), and I should now comment on their implications for the present work.

It has been stressed by many researchers that FOK judgments have measurable consequences as far as the initiation and termination of search, study time, and overall memory performance (e.g., Gruneberg et al., 1977; Koriat & Goldsmith, 1993b; Nelson & Narens, 1990). Thus, in terms of the model sketched in Figure 1, for example, not only does the amount of accessible information affect FOK judgments but the latter may also feed back into the search process, thereby affecting the amount of additional information retrieved. Although it might seem futile to tease apart the two types of effects, it may actually be quite important to do so.

Consider, for example, measures of response latency obtained in the course of a FOK experiment. The accessibility model predicts that FOK judgments should increase with decreasing response latency of retrieving partial or complete information. In contrast, as far as the consequences of FOK are concerned, the correlation between FOK and response latency should be positive: The higher the initial FOK judgments, the more time should one spend on a question before giving up (or until an answer is found). The evidence gathered so far discloses both types of relationships: In the present study, response latencies were negatively correlated with FOK (Experiment 3). In other studies, in contrast, higher FOK judgments were associated with longer search times (e.g., Gruneberg et al., 1977; Nelson et al., 1984, 1990; Reder 1987, 1988). Both effects may be clearly seen in the results of a recent study by Costermans et al. (1992): Mean FOK judgments were positively correlated with the time spent on a question before giving up, consistent with the assumption that FOK drives search continuation. At the same time, when a target was retrieved, confidence ratings in the correctness of the target increased with decreasing retrieval time, consistent with the presumed effect of ease of access on

subjective confidence. Note that this latter correlation was obtained for both correct recalls as well as commission errors. Very similar results were also found by Nelson et al. (1990).

The distinction between those effects that are due to the antecedents of FOK (and subjective confidence) and those that are due to its consequents seems to also underlie some of the differences observed between omission and commission errors. Nelson and his associates found a high positive correlation between FOK and latency of omission errors. In contrast, the correlation of FOK with latency of commission errors was nil or even slightly negative (Nelson et al., 1982, 1984, 1990). Although the former correlation apparently reflects the presumed effects of FOK on search time, the latter may represent a mixture of the two conflicting effects. This interpretation implies that FOK judgments elicited following retrieval of the complete target (correct or false) may represent a mixture of prospective and retrospective confidence judgments. This mixture may differ between individuals, as suggested by the very high interindividual variability in the size and direction of the correlation between FOK and latency of commission errors (see Nelson et al., 1984). If this interpretation is correct, FOK judgments following correct recall (normally not collected in FOK experiments) should evidence the same correlational pattern as that observed for commission errors.

The foregoing discussion suggests that the correlation between FOK and latency measures can disclose the extent to which a particular response property taps the antecedents of FOK or its consequents. Thus, because the results are primarily correlational, it might have been argued that, in fact, it is the FOK that is the cause for the greater accessibility of information in my experiments rather than vice versa. However, the correlational pattern between FOK, amount of information recalled, and recall latency is consistent with the assumption that it is the bases of FOK that are being tapped: Although there was a positive correlation between FOK and PI-T, each of them was negatively related to recall latency (see Figure 12).

In summary, the feedback loop between monitoring and memory search introduces certain complications that must be taken into account in specifying the extent to which FOK judgments obtained at time t_1 can predict memory performance at time t_2 .

In conclusion, I attempted to address in this article several issues concerning the FOK phenomenon within a common conceptual framework. It is my belief that the trace-access view, implicitly endorsed in many discussions, has greatly impeded serious efforts toward a deeper understanding of the internal machinery underlying the FOK. Specifically, it has steered research away from the experimental investigation of the determinants of FOK and the reasons for its predictive validity. This article made a modest effort toward the demystification of the FOK phenomenon and joins with some the current endeavors to gain insight into the question of how do we know that we know (e.g., Nelson & Narens, 1990; Reder & Ritter, 1992; Schwartz & Metcalfe, 1992).

References

- Begg, I., Duft, S., Lalonde, P., Melnick, R., & Sanvito, J. (1989). Memory predictions are based on ease of processing. *Journal of Memory and Language*, 28, 610-632.

- Blake, M. (1973). Prediction of recognition when recall fails: Exploring the feeling-of-knowing phenomenon. *Journal of Verbal Learning and Verbal Behavior*, 12, 311-319.
- Bousfield, W. A., & Rosner, S. R. (1970). Free vs. uninhibited recall. *Psychonomic Science*, 20, 75-76.
- Brown, A. S. (1991). A review of the tip-of-the-tongue experience. *Psychological Bulletin*, 109, 204-223.
- Brown, A. S., & Bradley, C. K. (1985). Semantic prime inhibition and memory monitoring. *Bulletin of the Psychonomic Society*, 23, 98-100.
- Brown, R., & McNeill, D. (1966). The "tip of the tongue" phenomenon. *Journal of Verbal Learning and Verbal Behavior*, 5, 325-337.
- Burke, D., MacKay, D. G., Worthley, J. S., & Wade, E. (1991). On the tip of the tongue: What causes word finding failures in young and older adults? *Journal of Memory and Language*, 30, 542-579.
- Carroll, M. L., & Nelson, T. O. (1993). Effect of overlearning on the feeling of knowing is more detectable in within-subject than in between-subject designs. *American Journal of Psychology*, 106, 227-235.
- Carroll, M., & Simington, A. (1986). The effects of degree of learning, meaning, and individual differences on the feeling of knowing. *Acta Psychologica*, 61, 3-16.
- Cohen, G. (1990). Why is it difficult to put names to faces? *British Journal of Psychology*, 81, 287-297.
- Costermans, J., Lories, G., & Ansay, C. (1992). Confidence level and feeling of knowing in question answering: The weight of inferential processes. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 18, 142-150.
- DaPolito, F., Guttenplan, H., & Steinitz, H. (1968). Accuracy of subjective judgments of information in long-term memory. *Psychonomic Science*, 13, 227-228.
- Durso, F. T., & Shore, W. J. (1991). Partial knowledge of word meanings. *Journal of Experimental Psychology: General*, 120, 190-202.
- Eysenck, M. W. (1979). The feeling of knowing a word's meaning. *British Journal of Psychology*, 70, 243-251.
- Fischhoff, B., Slovic, P., & Lichtenstein, S. (1977). Knowing with certainty: The appropriateness of extreme confidence. *Journal of Experimental Psychology: Human Perception and Performance*, 3, 552-564.
- Fodor, J. A. (1983). *The modularity of mind*. Cambridge, MA: MIT Press.
- Freedman, J. L., & Landauer, T. K. (1966). Retrieval of long-term memory: "Tip-of-the-tongue" phenomenon. *Psychonomic Science*, 4, 309-310.
- Gardiner, J. M., Craik, F. I. M., & Bleasdale, F. A. (1973). Retrieval difficulty and subsequent recall. *Memory & Cognition*, 7, 213-216.
- Glenberg, A. M., Sanocki, T., Epstein, W., & Morris, C. (1987). Enhancing calibration of comprehension. *Journal of Experimental Psychology: General*, 116, 119-136.
- Gruneberg, M. M. (1978). The feeling of knowing, memory blocks and memory aids. In M. M. Gruneberg & P. Morris (Eds.), *Aspects of memory* (pp. 187-209). London: Methuen.
- Gruneberg, M. M., & Monks, J. (1974). Feeling of knowing and cued recall. *Acta Psychologica*, 38, 257-265.
- Gruneberg, M. M., Monks, J., & Sykes, R. N. (1977). Some methodological problems with feeling of knowing studies. *Acta Psychologica*, 41, 365-311.
- Gruneberg, M. M., Smith, R. L., & Winfrow, P. (1973). An investigation into response blocking. *Acta Psychologica*, 37, 187-196.
- Gruneberg, M. M., & Sykes, R. N. (1978). Knowledge and retention: The feeling of knowing and reminiscence. In M. M. Gruneberg, P. E. Morris, & K. N. Sykes (Eds.), *Practical aspects of memory* (pp. 189-196). San Diego, CA: Academic Press.
- Hart, J. T. (1965). Memory and the feeling-of-knowing experience. *Journal of Educational Psychology*, 56, 208-216.
- Hart, J. T. (1966). Methodological note on feeling-of-knowing experiments. *Journal of Educational Psychology*, 57, 347-349.
- Hart, J. T. (1967a). Memory and the memory-monitoring process. *Journal of Verbal Learning and Verbal Behavior*, 6, 685-691.
- Hart, J. T. (1967b). Second-try recall, recognition and the memory-monitoring process. *Journal of Educational Psychology*, 58, 193-197.
- Jacoby, L. L., & Brooks, L. R. (1984). Nonanalytic cognition: Memory, perception, and concept learning. In G. H. Bower (Ed.), *The psychology of learning and motivation: Advances in research and theory* (Vol. 18, pp. 1-47). San Diego, CA: Academic Press.
- Jacoby, L. L., & Dallas, M. (1981). On the relationship between autobiographical memory and perceptual learning. *Journal of Experimental Psychology: General*, 3, 306-340.
- Jacoby, L. L., & Kelley, C. M. (1987). Unconscious influences of memory for a prior event. *Personality and Social Psychology Bulletin*, 13, 314-336.
- Jacoby, L. L., & Kelley, C. M. (1991). Unconscious influences of memory: Dissociations and automaticity. In D. Milner & M. Rugg (Eds.), *The neuropsychology of consciousness* (pp. 201-233). San Diego, CA: Academic Press.
- Jacoby, L. L., Kelley, C. M., & Dywan, J. (1989). Memory attributions. In H. L. Roediger & F. I. M. Craik (Eds.), *Varieties of memory and consciousness: Essays in honour of Endel Tulving* (pp. 391-422). Hillsdale, NJ: Erlbaum.
- Jacoby, L. L., Lindsay, S. D., & Toth, J. P. (1992). Unconscious influences revealed: Attention, awareness, and control. *American Psychologist*, 47, 802-809.
- Jacoby, L. L., & Whitehouse, K. (1989). An illusion of memory: False recognition influenced by unconscious perception. *Journal of Experimental Psychology: General*, 118, 126-135.
- James, W. (1893). *The principles of psychology* (Vol. I). New York: Holt.
- Jameson, K. A., Narens, L., Goldfarb, K., & Nelson, T. O. (1990). The influence of subthreshold priming on metamemory and recall. *Acta Psychologica*, 73, 55-68.
- Jones, G. V. (1989). Back to Woodworth: Role of interlopers in the tip-of-the-tongue phenomenon. *Memory & Cognition*, 17, 69-76.
- Jones, G. V., & Langford, S. (1987). Phonological blocking in the tip-of-the-tongue state. *Cognition*, 26, 115-122.
- Kelley, C. M., & Jacoby, L. L. (1990). The construction of subjective experience: Memory attributions. *Mind & Language*, 5, 49-68.
- Kolers, P. A., & Palef, S. R. (1976). Knowing not. *Memory & Cognition*, 4, 553-558.
- Koriat, A. (1975). Phonetic symbolism and feeling of knowing. *Memory & Cognition*, 3, 545-548.
- Koriat, A. (1976). Another look at the relationship between phonetic symbolism and the feeling of knowing. *Memory & Cognition*, 4, 244-248.
- Koriat, A., & Goldsmith, M. (1993a). *Metaphors for memory assessment: I. Comparing quantity-oriented and accuracy-oriented approaches*. Manuscript submitted for publication.
- Koriat, A., & Goldsmith, M. (1993b). *Metaphors for memory assessment: II. The strategic control of memory accuracy*. Manuscript submitted for publication.
- Koriat, A., Liechtenstein, S., & Fischhoff, B. (1980). Reasons for confidence. *Journal of Experimental Psychology: Human Learning and Memory*, 6, 107-118.
- Koriat, A., & Lieblich, I. (1974). What does a person in a "TOT" state know that a person in a "don't know" state doesn't know? *Memory & Cognition*, 2, 647-655.
- Koriat, A., & Lieblich, I. (1977). A study of memory pointers. *Acta Psychologica*, 41, 151-164.

- Krinsky, R., & Nelson, T. O. (1985). The feeling of knowing for different types of retrieval failure. *Ada Psychologies*, 58, 141-158.
- Lachman, J. L., Lachman, R., & Thronesbery, C. (1979). Metamemory through the adult life span. *Developmental Psychology*, 15, 543-551.
- Leonesio, R. J., & Nelson, T. O. (1990). Do different metamemory judgments tap the same underlying aspects of memory? *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 16, 464-470.
- Lewin, K. (1935), *A dynamic theory of personality*. New York: McGraw-Hill.
- Lichtenstein, S., & Fischhoff, B. (1977). Do those who know more also know more about how much they know? *Organizational Behavior and Human Performance*, 20, 159-183.
- Lichtenstein, S., Fischhoff, B., & Phillips, L. D. (1982). Calibration of probabilities: The state of the art to 1980. In D. Kahneman, P. Slovic, & A. Tversky (Eds.), *Judgment under uncertainty: Heuristics and biases* (pp. 306-334). Cambridge, England: Cambridge University Press.
- Lovelace, E. (1987). Attributes that come to mind in the TOT state. *Bulletin of the Psychonomic Society*, 25, 370-372.
- Lupker, S. J., Harbluk, J. L., & Patrick, A. S. (1991). Memory for things forgotten. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 17, 897-907.
- Mandler, G. (1991). Your face looks familiar but I can't remember your name: A review of dual process theory. In W. E. Hockley & S. Lewandowsky (Eds.), *Relating theory and data: Essays on human memory in honor of 'Beimel B. Murdoch* (pp. 207-225). Hillsdale, NJ: Erlbaum.
- McClelland, J. L. M., & Rumhart, D. E. (1985). Distributed memory and the representation of general and specific information. *Journal of Experimental Psychology: General*, 114, 159-188.
- Metcalf, J. (1986). Feeling of knowing in memory and problem solving. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 12, 288-294.
- Metcalf, J. (1993). Novelty monitoring, metacognition, and control in a composite holographic associative recall model: Implications for Korsakoff amnesia. *Psychological Review*, 100, 3-22.
- Metcalf, J., & Wiebe, D. (1987). Intuition in insight and noninsight problem solving. *Memory & Cognition*, 15, 238-246.
- Morris, C. C. (1990). Retrieval processes underlying confidence in comprehension judgments. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 16, 223-232.
- Mozur, M. C. (1991). *The perception of multiple objects: A connectionist approach*. Cambridge, MA: MIT Press.
- Nelson, T. O. (1984). A comparison of current measures of the accuracy of feeling-of-knowing predictions. *Psychological Bulletin*, 95, 109-133.
- Nelson, T. Q., & Dunlosky, J. (1991). When people's judgments of learning (JOLs) are extremely accurate at predicting subsequent recall: The "delayed JOL effect." *Psychological Science*, 2, 267-270.
- Nelson, T. O., Dunlosky, J., White, D. M., Steinberg, J., Townes, B. D., & Anderson, D. (1990). Cognition and metacognition at extreme altitudes on Mount Everest. *Journal of Experimental Psychology: General*, 119, 367-374.
- Nelson, T. Q., Gerler, D., & Narens, L. (1984). Accuracy of feeling-of-knowing judgment for predicting perceptual identification and relearning. *Journal of Experimental Psychology: General*, 113, 282-300.
- Nelson, T. Q., Leonesio, R. J., Landwehr, R. F., & Narens, L. (1986). A comparison of three predictors of an individual's memory performance: The individual's feeling of knowing versus the normative feeling of knowing versus base-rate item difficulty. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 12, 279-287.
- Nelson, T. O., Leonesio, R. J., Shimamura, A. P., Landwehr, R. F., & Narens, L. (1982). Overlearning and the feeling of knowing. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 8, 279-288.
- Nelson, T. O., & Narens, L. (1980). A new technique for investigating the feeling of knowing. *Ada Psychologies* 46, 69-80.
- Nelson, T. O., & Narens, L. (1990). Metamemory: A theoretical framework and new findings. In G. Bower (Ed.), *The psychology of learning and motivation, advances in research and theory* (Vol. 26, pp. 125 - 173). San Diego, CA: Academic Press.
- Nisbett, R. E., & Bellows, N. (1977). Verbal reports about causal influences on social judgments: Private access versus public theories. *Journal of Personality and Social Psychology*, 35, 613-624.
- Nisbett, R. E., & Wilson, T. D. (1977). Telling more than we know: Verbal reports on mental processes. *Psychological Review*, 84, 231-279.
- Read, J. D., & Bruce, D. (1982). Longitudinal tracking of difficult memory retrievals. *Cognitive Psychology*, 14, 280-300.
- Reason, J., & Lucas, D. (1984). Using cognitive diaries to investigate naturally occurring memory blocks. In J. E. Harris & P. E. Morris (Eds.), *Everyday memory, actions and absentmindedness* (pp. 53-69). San Diego, CA: Academic Press.
- Reder, L. M. (1987). Strategy selection in question answering. *Cognitive Psychology*, 19, 90-138.
- Reder, L. M. (1988). Strategic control of retrieval strategies. *The Psychology of Learning and Motivation*, 22, 227-259.
- Reder, L. M., & Ritter, F. E. (1992). What determines initial feeling of knowing? Familiarity with question terms, not with the answer. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 18, 435-451.
- Ryan, M. P., Petty, C. R., & Wenzlaff, R. M. (1982). Motivated remembering efforts during tip-of-the-tongue states. *Ada Psychologies* 51, 137-147.
- Schacter, D. L. (1983). Feeling of knowing in episodic memory. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 9, 39-54.
- Schacter, D. L., & Worling, J. R. (1985). Attribute information and the feeling of knowing. *Canadian Journal of Psychology*, 39, 467-475.
- Schwartz, B. L., & Metcalfe, J. (1992). Cue familiarity but not target retrievability enhances feeling-of-knowing judgments. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 18, 1074-1083.
- Senders, J. W., & Moray, N. P. (1991). *Human error: Cause, prediction, and reduction*. Hillsdale, NJ: Erlbaum.
- Stroop, J. R. (1935). Studies of interference in serial verbal reactions. *Journal of Experimental Psychology*, 18, 643-662.
- Tulving, E., & Madigan, S. A. (1970). Memory and verbal learning. *Annual Review of Psychology* 21, 434-437.
- Tulving, E., & Pearlstone, Z. (1966). Availability versus accessibility of information in memory for words. *Journal of Verbal Learning and Verbal Behavior*, 5, 381-391.
- Tversky, A., & Kahneman, D. (1973). Availability: A heuristic for judging frequency and probability. *Cognitive Psychology*, 4, 207-232.
- Van Bergen, A. (1968). *Task interruption*. Amsterdam: North-Holland.
- Whittlesea, B. W. A., Jacoby, L. L., & Girard, K. A. (1990). Illusions of immediate memory: Evidence of an attributional basis for feelings of familiarity and perceptual quality. *Journal of Memory and Language*, 29, 716-732.
- Wickelgren, W. A. (1970). Multitrac strength theory. In D. A. Norman (Ed.), *Models of human memory* (pp. 65-102), San Diego, CA: Academic Press.
- Yaniv, L., & Meyer, D. E. (1987). Activation and metacognition of inac-

cessible stored information: Potential bases for incubation effects in problem solving. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 13, 187-205.

Yaniv, I., Yates, J. F., & Smith, J. E. K. (1991). Measures of discrimination skill in probabilistic judgment. *Psychological Bulletin*, 110, 611-617.

Yavutz, H. S., & Bousfield, W. A. (1959). Recall of connotative meaning. *Psychological Reports*, 5, 319-320.

Received April 14, 1992

Revision received January 12, 1993

Accepted March 22, 1993 ■