1

INTRODUCTION AND SUMMARY

After a long period of time during which stimulus-response relations were at the focus of attention, research in psychology is now seeking to understand in detail the mechanisms and internal structure of cognitive processes that produce these relations. In the limiting case, we would like to have process models so explicit that they could actually produce the predicted behavior from the information in the stimulus.

This concern for the course of the cognitive processes has revived interest in finding ways to increase the temporal density of observations so as to reveal intermediate stages of the processes. Increasingly, investigators record the directions of the subject's gaze (eye movements), and the intermediate behaviors (movements or physical manipulations of stimulus material) that precede the solution or criterion performance. Since data on intermediate processing are costly to gather and analyze, it is important to consider carefully how such data can be interpreted validly, and what contribution they can make to our understanding of the phenomena under study.

One means frequently used to gain information about the course of the cognitive processes is to probe the subjects' internal states by verbal methods. These methods are the topic of this monograph.

USING VERBAL REPORTS: SOME ISSUES

There are several issues that we must deal with if we are to use subjects' reports as fundamental data in psychological experiments. First, we must respond to the strong doubts that have been expressed by many psychologists in the past about the suitability of subjects' verbalizations as scientific data. Second, we must consider the processing that must take

place in order to transform subjects' behaviors (whether verbal or not) into data. Third, we must examine how the encoding of behavior into data can be made objective and univocal, so that the resulting data will be "hard" and not "soft." Fourth, we must be explicit about the theoretical presuppositions that are necessarily embedded in the encoding process. Finally, we must specify the processes that allow us to go backward from the data to the behavior and thence to inferences about the subjects' thought processes.

We offer a few comments on each of these five issues. They will reappear frequently as recurrent themes throughout the monograph.

Doubts About Verbal Data

Since the triumph of behaviorism over "introspectively" oriented competing viewpoints, verbal reports have been suspect as data. More precisely, behaviorism and allied schools of thought have been schizophrenic about the status of verbalizations as data. On the one hand, verbal responses (or key punches that are psychologically indistinguishable from verbal responses, except that they are made with the finger instead of the mouth) provide the basic data in standard experimental paradigms. In a concept attainment experiment, the subjects say (or signal) "yes" or "no" when a possible instance is presented to them. In a problem solving experiment, they report the answer when they find it. In a rote verbal learning experiment, they say "DAX" when the stimulus syllable "CEF" is presented. The actual performance measures commonly used–latencies and numbers of items correct– are derived from these responses, and the former depend for their validity on the veridicality of the latter.

On the other hand, modern psychology has been dubious about verbalizations produced by subjects along the route to their solutions or final responses. Even more dubious has been the status of responses to experimenter probes or retrospective answers to questions about prior behavior. All of these sorts of verbal behavior are frequently dismissed as variants of the discredited process of introspection (Nisbett & Wilson, 1977). Introspection, it has generally been argued, may be useful for the discovery of psychological processes; it is worthless for verification. As Lashley (1923, p. 352) said, in a vigorous and widely cited attack on the method, "introspection may make the preliminary survey, but it must be followed by the chain and transit of objective measurement."

Extracting Data from Behavior

The notion that verbal reports provide possibly interesting but only informal information, to be verified by other data, has affected the ways in which verbalizations are collected and analyzed. If the purpose of obtaining verbal reports is mainly to generate hypotheses and ideas, investigators need not concern themselves (and generally have not concerned themselves) with methodological questions about data collection. As a result, there is little published literature on such issues, the data-gathering and data-analysis methods actually used vary tremendously, and the details of these methods are reported sketchily in research publications that make use of such data.

If we are to make rapid and continuing progress in understanding human cognitive processes, this state of affairs is wholly unsatisfactory. In the first place, no clear guidelines are provided to distinguish illegitimate "introspection" from the many forms of verbal output that are routinely treated as data–as passing the chain and transit test (see the examples above). On what theoretical or practical grounds do we distinguish between the subject's "yes" or "no" in a concept attainment experiment and his assertion that the hypothesis he is entertaining is "small yellow circle"? In the second place, no distinctions are made among such diverse forms of verbalization as thinking-aloud (TA) protocols, retrospective responses to specific probes, and the classical introspective reports of trained observers. All are jointly and loosely condemned as "introspection."

Soft versus Hard Data

Some investigators call verbal reports and verbal descriptions "soft data" in contrast to simple behavioral measures like latency or correctness of response, which are referred to as "hard." What does this distinction mean? In science one would like to maintain as clear a separation as possible between data and theory. Data are supposed to derive directly from observation; theories are supposed to account for, explain, and predict these observation-based data. Data are "hard" when there is intersubjective agreement that they correspond to the facts of the observed behavior.

Even psychoanalytically or existentially oriented psychologists will

accept response latencies as data–even though being possibly irrelevant data for explaining behavior. When, however, an analyst codes a five-second description of a dream as "oral fixation," many psychologists would argue that this encoding is not a datum but a subjective interpretation of the data (i.e., of the verbal description of the dream). Surely, theory-laden inferences were required to derive the encoding from the verbal protocol. Data are regarded as "soft" to the degree that they incorporate such inferences, especially when the theoretical premises and rules of inference are themselves not completely explicit and objective. The problem with "soft" data is that different interpreters making different inferences will not agree in their encodings, and each interpreter is likely, wittingly or not, to arrive at an interpretation that is favorable to his theoretical orientation.

The hard-soft distinction is orthogonal to the distinction between verbal and non-verbal. The same problems of inference can emerge in observers' attempts to understand non-verbal events (e.g., sequences of physical movements, pieces of music). Such events may require as much interpretation as is required to understand verbal sequences.

Technological advances have enhanced our ability to treat verbal protocols as hard data. Until tape recorders were generally available, it was common practice for experimenters to take selective notes of verbalizations, paraphrasing and omitting whatever was "unimportant." In analyzing such notes further, it was impossible to distinguish the inferences from the original verbalizations. Using encodings of verbal protocols as data has often been made even more difficult because the theories employed, explicitly or implicitly, in the encoding were formulated in very general terms. The search for general mechanisms also led to overall interpretations of entire protocols with little concern for encoding and explicating individual protocol statements.

More recent research based on explicit information processing models of the cognitive process has caused thinking-aloud verbalizations to be viewed in a new light. It is now standard procedure to make careful verbatim transcripts of the recorded tapes, thus preserving the raw data in as "hard" a form as could be wished. At the same time, information processing models of the cognitive processes provide a basis for making the encoding process explicit and objective, so that the theoretical presuppositions entering into that process can be examined objectively.

Theoretical Presupposition in Encoding

Clyde Coombs, in his book *A Theory of Data*, shows that raw data go through a typical sequence of steps on the route from initial observation to the edited and encoded form in which they are used to test theories or make predictions. These steps, which are not neutral with respect to theory, can be seen in the processing of protocol data as they can with other kinds of data. At the first step, theory delimits a small portion of the universe of potentially observable behavior as being relevant. This judgment of relevance determines what behaviors should be recorded. At the next step, these behaviors are encoded in a manner that is again determined on theoretical grounds.

In the case of verbal behavior, the process begins with tape-recording, containing essentially all the auditory events that occurred during the experimental session. In producing from the tape a written transcript, some selection is required. After the temporal information, repetitions, and stress have been used to segment and parse the verbal stream, most of this information is usually eliminated from the transcript, except as it is captured by punctuation. We will refer to this transcription step as *preprocessing*.

At the next step, the preprocessed segments are encoded into the terminology of the theoretical model. This is often achieved by first determining coding categories, a priori, and then having human judges make the coding assessments. If each of the segments is to be treated as an independent datum, then the encoding of that segment must be made on the basis of the information contained in it, independently of the surrounding segments. In Chapter 6 of this book, we will discuss at some length methods for carrying out this kind of local encoding, and the conditions that must be met to make it possible.

Verbal protocols have been analyzed in two rather different ways. One method claims *not* to require the analysis of meanings, while the other does require it. In the first kind of analysis, subject and experimenter have agreed, by prior instruction, upon specific signals, which may be speech signals or button presses, for their communication. These signals are mostly arbitrary–a subject could say "cef" instead of "yes"; communication is possible only because of the agreement established between subject and experimenter. To analyze the recorded verbalizations under these conditions, the experimenter has only to categorize each speech signal into one of the agreed-upon categories. In theory, if not in practice, a coder should not even need to know the subject's

language–assuring that no meaningful analysis of inferencing is involved. A large number of paradigms in psychology use this kind of analysis. For example, studies using scales and multiple-choice alternatives can all be seen as instances of this method.

In the second kind of analysis, the observed verbalizations are analyzed in terms of their meanings. Even in this case, the theory buiding the analysis limits the encoding to selected aspects and features rather than the full meaning of the verbalization. For example, in a typical concept attainment task, each instance or stimulus can be represented as a unique combination of features. Each distinct concept can be represented by some particular configuration of features. Then encoding simply requires the mapping of the verbalizations onto these concepts and features usually a rather unequivocal matter. Although the space of logically possible different concepts may be very large, it is severely limited compared with the variability of natural language. Thus a verbalization like "red circles are cef's" can normally be encoded as identical with "blood-colored round ones are cef's."

The context of a particular theory and experiment greatly constrains the range of possible interpretation and allows the meaningful analysis of verbalizations to be selective and incomplete. If a theory of concept attainment is limited to the language of hypotheses, many verbalizations will not be encoded at all–statements like, "I wonder what I should do. I'll just guess on this one." Many examples can be cited of this kind of meaningful analysis, where verbalizations are mapped onto a priori formal alternatives. The analysis of memory for meaningful text has been studied by Kintsch (1974) and many others. Newell and Simon (1972) analyzed tasks, identifying formally defined knowledge states in terms of which subjects' thinking-aloud protocols could be encoded.

Many analyses of verbalizations do *not* fit the above scheme, including most analyses that seek to arrive at an understanding of the verbalizations. In less formal kinds of analysis, the encoding scheme is not defined formally and a priori, but the search for interpretations proceeds in parallel with the search for an appropriate model or theory. We recognize clearly the need for and value of such interactive processes in the search for theories in new domains, but in our own account here we will be concerned primarily with situations where the theoretical terms are fixed before the actual encoding begins.

Inferring Thought Processes From Behavior

It is sometimes believed that using verbal data implies accepting the subjects' interpretation of them or of the events that are reported. This issue of trust has its origins in our everyday experience and use of language. In order to communicate effectively with other people, we accept their word for many facts. If someone says that he has bought a new car, we generally accept his statement as true instead of asking him to produce the sales contract or a receipt. In a similar vein we trust people–at least our friends–to answer questions correctly and to give us the best advice they can. However, if the issue is important to us or we suspect ulterior motives in the responses, we may demand more details and may review all the available evidence ourselves. The same thing holds in scientific research; few scientists will accept another scientist's claim of finding conclusive evidence for ESP without wanting an independent review of the evidence.

Subjects' reports of their own mental states and mental processes raise slightly different issues of trust. According to a naive theory of consciousness, subjects have the sole *direct* access to their own mental states and processes. The subjective feeling of one's ability to report one's own mental experiences veridically is strong. For a great many reasons, this confidence is not shared by experimental psychologists, who have shown that under numerous circumstances such self-reports are unreliable.

However, the issue of the reliability of self-reports can (and, we think, should) be avoided entirely. The report "X" need not be used to infer that X is true, but only that the subject was able to say "X"–(i.e., had the information that enabled him to say "X.") By following this path, we can even show that there is an inverse relation between how much subjects need to be trusted and how much information they verbalize. For the more information conveyed in their responses, the more difficult it becomes to construct a model that will produce precisely those responses adventitiously–hence the more confidence we can place in a model that does predict them.

Consider, for example, the following possible interchanges between experimenter and subject:

1. Do you know the name of the capital of Sweden? *Yes.*
2. Which of these three, Oslo, Stockholm, or Copenhagen, is the capital of Sweden? *Stockholm.*
3. Name the capital of Sweden. *Stockholm.*

4. (A retrospective report as to how the subject arrived at an answer to Question 1): *First I tried to picture where Sweden is located on a map of Europe, then Oslo came to mind, but I remembered that it is the capital of Norway. Then Stockholm popped up and I remembered that is where the Nobel prizes are awarded; then I felt sure I could answer "yes."*

In the first case we have to trust the subject if we want to infer that he actually knows the capital, whereas in the third case it is unlikely that he could generate the correct name unless it were accessible from memory. The primary difference between second and third cases is that, for the second, one could conceive of a number of processes other than memory retrieval (e.g., guessing) that would account for the response. The fourth response, the retrospective report, also verifies that the subject has the name in memory together with some redundant information about it that gives him confidence in his answer. Of course we do not have to believe that he has given a veridical report of the process whereby he generated the name, although there is nothing implausible about the sequence of associations he reports.

Consider next a more controversial example, which has played a role in the psychological literature on learning without awareness. After a learning experiment, the experimenter asks the subjects whether they were aware of any relation between the stimuli and responses, on the one hand, and the reward contingencies on the other. Yes/no responses to this question are informative only if we trust the subjects. If a subject, however, describes the stimulus-response contingency for reward, we can be reasonably certain that he had access to this information while he was learning. On the other hand, if a subject is unable to report anything about the contingency, we *cannot* conclude that he wasn't aware of it during the learning process–we have solid evidence neither for nor against awareness during the experiment. Later, we will discuss the problem of making inferences from reports of lack of information.

These examples illustrate that the information externalized in verbal responses often provides the experimenter with data that eliminate the need for trust in the subject. The examples also show that verbal reports may be generated in many ways. To understand the reports, we must understand the processes by which they were generated. In none of these respects do data from verbal reports differ from data based on other types of observations.

Some Basic Assumptions

We can now summarize the basic assumptions that set the stage for our further explorations. Most fundamentally, we see verbal behavior as one type of recordable behavior, which should be observed and analyzed like any other behavior. The cognitive processes that generate verbalizations are a subset of the cognitive processes that generate any kind of recordable response or behavior. Hence, we would look for the same kind of "mechanical" and complete process description of verbal behavior as of other kinds of behavior, and we would not accept magical or privileged processes as explanations for verbalizations.

Whether one can and should trust subjects' verbal reports is not a matter of faith but an empirical issue on a par with the issue of validating other types of behavior, like eye fixations or motor behavior. A single invalid verbal report should not force us to discard analysis of verbal reports generally. Indeed, this monograph will undertake to build a theory of verbalization, so that we can then specify when, where, and under what kinds of instructions informative verbal reports can be obtained from subjects.

Postulating that the cognitive processes underlying verbalization are a subset of all cognitive processes implies that verbalization must comply with the constraints that have been identified, experimentally, to govern all cognitive processes. These information processing constraints will provide powerful guidelines for our attempts to specify how observed verbalizations could have been generated. We wish to account for verbally reported information by proposing a processing model sufficiently powerful to regenerate that information.

Plan of Attack

Our first task is to describe a general theory of cognitive processes and structure, which, we argue, accounts for verbalizations and verbal reports. For reasons that have already been stated, the analysis must be carried out within a framework of theory. This framework must be sufficiently general to permit us to relate, within a unified perspective, all the kinds of data that are commonly used in psychological experiments.

Usually, in choosing between theories, we want to pick the strongest one–the one that will make the strongest predictions. In the

present case, where the theory we choose will influence the way in which we encode and analyze our data, we want to pick the weakest and most "neutral" one that can do the job. The fewer controversial assumptions we incorporate in the theory, the less we will be involved in the circularity of using theory-laden data to test our theories. Nevertheless, there appears to be no way of processing data that does not incorporate *some* theoretical assumptions about the system and processes that generated the data. Our particular strategy will be to set forth the theory in its most general, hence least controversial, form first, then add more specific hypotheses where they are required.

After presenting the theory as an information processing model of cognitive processes, we will survey the literature on verbal reporting and derive from it a taxonomy of reporting procedures. We will follow this survey with an historical review of earlier approaches to verbal reports. We will then take up the major issues surrounding the use and validity of verbal reports, discussing the empirical studies within the framework of a more detailed information processing model.

THE PROCESSING MODEL

Our purpose in presenting a specific processing model is to aid us in interpreting verbal data obtained from subjects and the relation of their verbal to their other behavior. Since the data (including the verbal data) are gathered in order to test theories about the human information processing system, we are engaged in something of a bootstrap operation. We need a model in order to interpret data that are to be used, in turn, to test the model. Under these circumstances, our data-interpretation model should be as simple as possible, and it must not incorporate components that are themselves bones of theoretical contention. The model should be robust (i.e., compatible with a wide range of alternative assumptions about human information processing).

The specifications we are about to present are simple and robust in this sense, and, indeed, summarize the core that is common to most current information processing theories of cognition. Of course they are not entirely neutral, for they would be hard to reconcile with an extreme form of behaviorism that denied the relevance of central processes to the explanation of behavior. But they are not specific to the view of any particular "sect" within the general information-processing tradition. (For fuller discussion of the model, see Newell and Simon (1972, Chapter 14), and Simon (1979, Chapters 2, 3).

General Specification

The most general and weakest hypothesis we require is that human cognition is information processing: that a cognitive process can be seen as a sequence of internal states successively transformed by a series of information processes. An important, and more specific, assumption is that information is stored in several memories having different capacities and accessing characteristics: several sensory stores of very short duration, a short-term memory (STM) with limited capacity and/or intermediate duration, and a long-term memory (LTM) with very large capacity and relatively permanent storage, but with slow fixation and access times compared with the other memories.

Within the framework of this information processing model, it is assumed that information recently acquired (attended to or heeded)* by the central processor is kept in STM, and is directly accessible for further processing (e.g., for producing verbal reports), whereas information from LTM must first be retrieved (transfered to STM) before it can be reported.

This general picture is compatible with all sorts of specific hypotheses that have been put forth with respect to the details of the mechanisms. For example, some theorists propose that what we call "short-term memory" is not a separate, specialized store but simply a portion of LTM that is currently and temporarily activated (Anderson, 1976). Some theorists believe that information in STM extinguishes with passage of time, unless rehearsed; others that it is lost only when replaced. In general, these differences of detail do not affect the model at the level of specificity required for our purposes. The important hypothesis for us is that, due to the limited capacity of STM, only the most recently heeded information is accessible directly. However, a portion of the contents of STM are fixated in LTM before being lost from STM, and this portion can, at later points in time, sometimes be retrieved from LTM.

Our specification of the system is general, but it is not vague. Specific information processing models that incorporate these features have been constructed in the form of computer programs, and these have

*Because the phrase "attended to" is often stylistically awkward, we will sometimes use "heeded" instead. So we will say, more or less synonymously, that information was "attended to," was "heeded," or was "stored in STM."

been shown to produce a variety of behaviors previously observed in psychological laboratories. Verbal predictions of how such a system behaves can, thereby, be tested by using a computer program as a simulator. The principal model of this kind that guides our own thinking about these processes is the EPAM program, due to Feigenbaum (1963) and Simon, and discussed in some detail in Section 3 of Simon (1979).

We assume that any verbalization or verbal report of the cognitive processes would have to be based on a subset of the information held in STM and LTM. From this and the above hypotheses, the taxonomy of verbalization procedures shown in Table 1-1 follows in a straightforward fashion (Ericsson & Simon, 1980).

Table 1-1

A Classification of Different Types of Verbalization Procedures as a Function of Time of Verbalization (Rows) and the Mapping From Heeded to Verbalized Information (Columns)

	Relation between heeded and verbalized information			
		Intermediate processing		
Time of verbalization	Direct one to one	Many to one	Unclear	No relation
While information is attended	Talk aloud Think aloud			
While information is still in short-term memory	Concurrent probing	Intermediate inference and generative processes		
After the completion of the task-directed processes	Retrospective probing	Requests for general reports	Probing hypothetical states	Probing general states

The two dimensions of Table 1-1 represent two major distinctions. First, the time of verbalization is important in determining from what memory the information is likely to be drawn. Second, we make a distinction between procedures where the verbalization is a direct articulation or explication of the stored information, and procedures where the stored information is input to intermediate processes, like abstraction and inference, so that the verbalization is a product of this intermediate processing.

Detailed Specification

We now specify more fully the components of the information processing system that we have just sketched. The model draws upon a variety of sources that are summarized in Newell and Simon (1972, Ch. 14) and Simon (1979, Ch. 2.3). Few of the model's specifications are controversial. It makes no real difference, for example, whether we assume a single homogeneous memory with different modes of activation (e.g., Anderson, 1976; Shiffrin & Schneider, 1977) or several discrete memory stores (sensory stores, STM, and LTM). The important matters, which can be described in either terms, relate to the amounts and kinds of information that can be retained, and the conditions for accessing them and reporting them verbally. We will use the conventional model of multiple memories in our description.

Recognition. Information received from the sensory organs resides for a short time in memories (iconic and echoic memories) associated with the different senses. During this time, portions of the sensory information are directly *recognized* and encoded with the aid of information already stored in LTM. Recognition associates the stimulus, or some part of it, with existing patterns in LTM, and stores in STM "pointers" to those familiar patterns. (The EPAM discrimination net is a model of this recognition mechanism.) Intermediate stages of the direct recognition process (the successive steps of discrimination), which may take only 10 to 100 msec, do not use STM to store their products.

Long-Term Memory. The LTM may be pictured as an enormous collection of interrelated nodes. Nodes can be accessed either by recognition (through the discrimination net), as just explained, or by way of links that associate these nodes to others that have already been accessed. Information accessed in either way is then represented by pointers in STM. Thus, information can be brought into STM from sensory stimuli via the recognition process, or from LTM via the association process. Association processes are much slower than direct recognition processes, requiring at least several hundred msec for each associative step. Associative processes may use STM to store intermediate steps. So, for example, in recalling a name that is not immediately accessible, a person may use a sequence of cues to find an associative path, step by step, to the sought-for name. Such processes may last tens of seconds, or even minutes, and may leave numerous intermediate symbols in STM, where they are temporarily available for verbal reports.

Short-term Memory. The central processor (CP), which controls and regulates the non-automatic cognitive processes, determines what small part of the information in sensory stimuli and LTM finds its way into STM. This is the information that is *heeded* or *attended to*. The amount of information that can reside in STM at one time is limited to a small number (four?) of familiar patterns (*chunks*). Each chunk is represented by one symbol or pointer to information in LTM (Simon 1979, Ch. 2.2). As new information is heeded, information previously stored in STM may be lost.

When a cognitive task (e.g., mental addition of a column of figures) is being carried out, the typical chunks in STM are pointers to the operands, operators, and outputs of the operations that are being performed. Thus, in adding 3 to 4, pointers corresponding to the symbols "3," "4," "PLUS," and "7" might at some time be present in STM. Since, in our culture, adding two digits involves a direct reference to LTM ("table lookup"), no further detail of the process would be heeded in STM or available for verbal reports. On the other hand, if the task were to multiply 17 by 45, STM might hold, at various points in the process "45," "17," "7," "TIMES," "3" (the carry in multiplying 45 by 7), "315" (the first intermediate product), "45," "1," "TIMES," "PLUS," "765."

We hold no brief for the details of the above description, which is intended merely as an example of the *kinds* of information we would expect to be heeded in STM, and to be available, potentially, for concurrent or retrospective reports. The specific details would depend on the particular strategies subjects used and the nature of the chunks they had stored in LTM (Simon, 1979, Ch. 2.4). STM would symbolize the process only down to some modest level of detail (corresponding to elementary processes of a second or two in duration), and we would not expect to find information there about simple, automated processes (e.g., the processes of retrieval from LTM or recognition processes), much less about neuronal events. Thus, the architecture of the control apparatus (CP) determines the fineness of grain of the representation of processes in STM.

Control of Attention. The flow of attention is diverted, from time to time, by interruptions through the higher control mechanism. Intermediate stages in these interruptions, not being symbolized in STM, are not reportable. Sudden movements in peripheral vision, loud noises, emotions operating through the reticular system are important causes of interruption and shift in attention (Simon, 1979, Ch. 1.3). While information heeded immediately before or after a shift in attention may sometimes allow subjects to give a relatively clear account of the interruption, we would expect such information to be less complete than reports of an orderly process that is induced by the successive content of STM itself (e.g., a thought sequence during which goals in STM are guiding the thought processes).

Fixation. New information is retained in STM during the time the CP is attending to it. In order to create an LTM representation of new information that can later be recalled, associations must be built up by coding and imaging, as well as new tests and branches in the recognition network. These learning processes, including the storage of new information in LTM and the addition of new pathways in the discrimination net for accessing it, are modeled in some detail by EPAM (Simon, 1979, Section 3). Processing of the order of 8 to 10 seconds is required to assemble each new chunk from its familiar components in STM, and to store it in LTM as a new chunk (Simon, 1979, Chs. 2.2, 2.3).

Automation. As particular processes become highly practiced, they become more and more fully automated. (Shiffrin & Schneider, 1977). Automation means that intermediate steps are carried out without being interpreted, and without their inputs and outputs using STM. The automation of performance is therefore quite analogous to executing a computer algorithm in compiled instead of interpretive mode. Automation (and compiling) have two important consequences. They greatly speed up the process (typically, by an order of magnitude) and they make the intermediate products unavailable to STM, hence unavailable also for verbal reports.

TYPES OF VERBALIZING PROCEDURES

The only feature common to the whole range of techniques used to obtain verbal data is that the subject responds orally to an instruction or probe. Because of the flexibility of language, there are virtually no limits to the probes we can insert and the questions we can ask subjects that will elicit some kind of verbal response.

Within our theoretical framework, we can represent verbal reporting

as bringing information into attention, then, when necessary, converting it into verbalizable code, and finally, vocalizing it. The crucial issue for verbal reporting procedures is what information is heeded. There have been studies showing that the response modality does not affect the frequency of different responses. Newhall and Roderick (1936) found no differences in frequencies between verbal reports, button presses with fingers, or pedal presses with the feet. This result indicates that the response is heeded symbolically, and then translated into the appropriate overt form. (See Chapter 5 for further discussion.)

Two forms of verbal reports can claim to being the closest reflection of the cognitive processes. Foremost are *concurrent verbal reports*–"talk aloud" and "think aloud" reports–where the cognitive processes, described as successive states of heeded information, are verbalized directly (see Figure 1-1).

We claim that cognitive processes are not modified by these verbal reports, and that task-directed cognitive processes determine what information is heeded and verbalized. We will evaluate this claim empirically in Chapter 2.

A second type of verbal report is the *retrospective report.* A durable (if partial) memory trace is laid down of the information heeded successively while completing a task. Just after the task is finished, this trace can be accessed from STM, at least in part, or retrieved from LTM and verbalized. Retrospective reports based on information in LTM will require an additional process of retrieval that will display some of the same kinds of error and incompleteness that are familiar from experimental research on memory. Both of these kinds of reports, we claim, are direct verbalizations of specific cognitive processes.

Recoding Before Verbalization

Various processes, and especially recoding processes, may intervene between the time information was heeded by the central processor (CP) and the time a verbalization is generated. When information is reproduced in the form in which it was heeded, we will speak of *direct* or *Level 1* verbalization. When one or more mediating processes occurs between attention to the information and its delivery, we will speak of *encoded* or *Level 2* or *Level 3* verbalization. A number of different kinds of intermediate processes between access and verbalization may modify the information. Among the important kinds are the following:

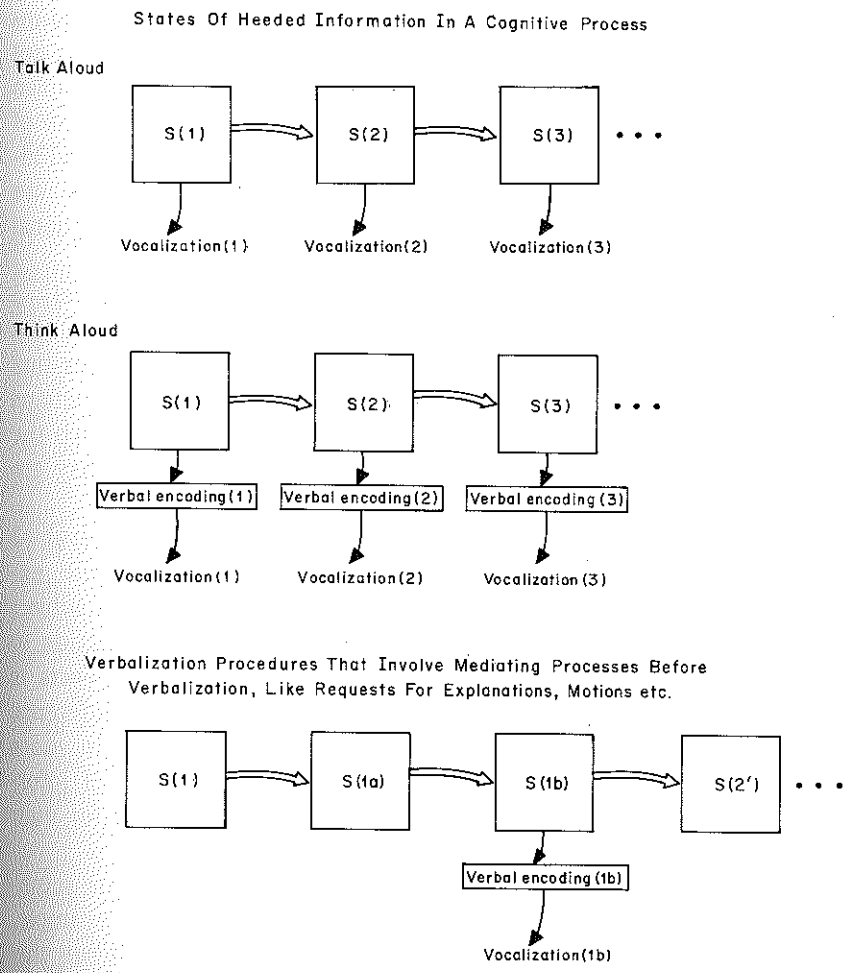

Figure 1-1

The Relation Between the Heeded States of a Cognitive Process and Verbal Reports for Various Types of Verbal Report Procedures

1. Recoding into verbal code (Level 2 verbalization). When the internal representation in which the information is originally encoded is not a verbal code, it has to be translated into that form. Werner and Kaplan (1963) have shown that when subjects generate verbal descriptions of nonverbal stimuli for their own future use, the format is compact and incorporates many idiosyncratic referents. When verbalizations are generated to communicate information to another person, additional processing is required to find referents (Werner & Kaplan, 1963).

2. Intermediate scanning or filtering processes (Level 3 verbalization). When the task instructions ask for verbalization of only selected information, it is necessary to postulate additional processes that test if the heeded information is of the desired type. Such instructions are used, for example, in commentary driving experiments, in which the subjects are asked to report all perceived traffic hazards while they are driving a car (Soliday & Allen, 1972).

3. Intermediate inference or generative processes (Level 3 verbalization) The situation is even more complicated if the experimenter is interested in particular aspects of the situation that a subject would not ordinarily attend to. The issue of whether the instruction to verbalize calls for information not normally heeded by the subjects is central and directly related to the occurrence of intermediate inference and generative processes. Since we will return to this issue in more depth, only a brief summary will be given here of the types of information that are likely to require additional mediating processing for their generation.

In addition to verbalizing their ongoing thinking, subjects are sometimes asked for verbal descriptions of their motor activities, for example, what objects are moved where, or where they are looking. When this information is not heeded directly, as is often the case, the subject is required to observe his or her own internal processes or overt behavior to generate the information.

Experimenters are often interested in subjects' reasons for their overt behavior and consequently ask the subjects to verbalize their motives and reasons, which may not be available directly or even at all. In an excellent review of research on the effects of persuasive messages, Wright (1980) discusses a wide range of biases due to different verbal report procedures.

In sum, with Level 1 and Level 2 verbalization the sequence of heeded information remains intact and no additional information is heeded. On the other hand Level 3 verbalization requires attention to additional information and hence changes the sequence of heeded information.

Retrospective Reports

In the ideal case the retrospective report is given by the subject immediately after the task is completed while much information is still in STM and can be directly reported or used as retrieval cues. It is clear that some additional cognitive processing is required to ascertain that the particular memory structures of interest are heeded. Our model predicts that retrospective reports on the immediately preceding cognitive activity can be accessed and specified without the experimenter having to provide the subject with specific information about what to retrieve. In this particular case, the subject will still retain the necessary retrieval cues in STM when a general instruction is given "to report everything you can remember about your thoughts during the last problem." This form of retrospective verbal report should give us the closest approximation to the actual memory structures.

Even in this favorable case, some problems arise that are common to all kinds of verbal reports from LTM. First, the retrieval operation is fallible, in that other similar memory structures may be accessed instead of those created by the just-finished cognitive process. The probability of this occurring increases markedly if the subjects have just solved a series of similar problems. However, since most accessed memory structures contain redundant information beyond the cues used for retrieval, subjects may use this additional information to validate the retrieval as well as to increase their confidence in the veridicality of the retrieved information. In a subsequent section we will discuss this type of evaluation further and examine the relevant theoretical and empirical literature.

A second general problem when retrieving cognitive structures is to separate information that was heeded at the time of a specific episode from information acquired previously or subsequently that is associated with it (Mueller, 1911). For example, if a picture reminds one of an old friend, it may be tempting to use the stored information about that friend to *infer* what the person in the picture looked like. (In Chapter 3 we will discuss this issue in more detail.) It may be possible to eliminate this artifact by instructing subjects only to report details that they can remember heeding at the time of the original episode (Mueller, 1911). By imposing a requirement of determinable memory as a basis for reporting, we can

avoid many subjects' tendency to fill in information that they can't remember but "must" have thought.

Inferential or Generative Processes

The most marked difference between concurrent and retrospective reporting is that retrospective reports refer to a cognitive process that is completed and cannot be altered and influenced. Hence, if subjects are requested to report information that was never heeded, they cannot possibly base their responses on direct memory. The subjects can answer that they don't know, but often they will infer and generate an answer on the basis of information provided in the question and other information accessible from LTM. Since retrieval from LTM may be an onerous task, even in situations where the information is potentially retrievable subjects may prefer to generate the information instead.

The most common probe that creates this problem is the why-question: for example, "Why did you do this?" or "Why did you prefer that product?" In an interesting discussion, Lazarsfeld (1935) points to many issues and problems in interpreting responses to why-questions, where subjects select one alternative out of several possibilities. Some of the alternatives may never have been heeded. If we wish to find out: "Why did you buy this book?" we may receive, out of the same concrete experience of the respondent, quite different answers, according to whether we stress "buy," "this," or "book." "If the respondent understood: 'Why did you BUY this book?' he might answer, 'Because the waiting list in the library was so long that I shouldn't have got it for two months.' If he understood: 'Why did you buy THIS book?' he might tell what interested him especially in the author. And if he understood: 'Why did you buy this BOOK?' he might report that he at first thought of buying a concert ticket with the money, but later realized that a book is a much more durable thing than a concert, and such reasoning caused him to spend his money upon the book" (Lazarsfeld, 1935, p. 29).

The example is instructive in showing that a person who did not actually buy the book, and hence had no specific memory of the associated cognitive processes, could give the same or similar answers as *plausible* reasons for someone else's buying a book. Hence, the answers can be generated (inferred) without access to a specific memory trace of the episode.

Directed or Specialized Probing

Verbal probes differ in the comprehensiveness of the topics to be reported and the generality or particularity of the events to be reported. Let us first consider topic specificity. In many studies, the investigator is interested only in particular aspects of subjects' behaviors. Then the verbal probe may be constructed to induce the subjects to generate information specifically relevant to the hypotheses under consideration. In order to help subjects retrieve the desired information from memory and to induce greater completeness of the verbal reports, the question or verbal probe often contains contextual information. To guard against subjectivity in analyzing verbal reports, the investigator often supplies subjects with a fixed set of alternative responses. In contrast, a general instruction to give verbal reports typically asks subjects to tell everything they can remember or are thinking of while performing the task.

In most cases, verbosity and absence of selectivity in subjects' reports is not an important problem. What the subject reports is likely to be less, rather than more, than we should like to hear. In no study known to us using general instructions has the investigator complained that subjects have reported too much information from actual memory.

One common difficulty in probing for specific information, especially when the subjects are offered a fixed set of alternative answers, is to know that the questions conform to the internal representations the subjects are employing in their thought. Probes for types of information that subjects don't have directly accessible, or probes that provide inadequate sets of alternatives may force subjects to intermediate and inferential processing, and hence produce verbal reports that are not closely related to the actual thought process. Moreover, when specific, fixed-alternative probes are used, there is no way to detect from subjects' responses that this has occurred.

Since providing contextual information and prompts to subjects may aid recall from LTM, in studies of LTM the use of prompts and context is frequent and relatively well-motivated. When subjects are asked to report on immediately preceding cognitive processes of relatively short duration, specific probes are more questionable and less useful. In a logical sense, the experimenter gets just as much information from the subject in the third as in the first two of the following three cases.

(1) Directed probe 1
Question: Did you use X as a subgoal?
Answer: Yes.

(2) Directed probe 2
Question: Did you use any subgoals? If so, which?
Answer: Yes, I used X.

(3) Undirected probe
Verbal report: ...I was first trying to get X and I...
when I attained X...

The replies in all three cases provide evidence that the subject used X as a subgoal, yet the evidence is stronger in the third case than in the second, and in the second than in the first. The verbalization in the first case could easily be generated by processes independent of any memory for the actual thought processes. Comparing the second and third cases, the former communicates to subjects what information the experimenter expects them to report. It may encourage subjects to try to infer or guess what particular information the experimenter will accept, and to generate information accordingly.

In many cases, other criteria are available for estimating the validity of the reports. An analysis of the task (Newell & Simon, 1972) will often provide strong indications of the adequacy of verbalized information, especially in cases with many logical possibilities for response.

Finally, different kinds of probes may have different effects upon the behavior of subjects. Requesting a certain kind of information may suggest to subjects what aspects of the task are important. Subjects may also alter their normal ways of processing so as to be able to give the requested information to the experimenter on subsequent trials.

In studies that use retrospective verbalization, subjects are seldom asked what they can remember about specific instances of their cognitive processes. Rather, they are usually asked to retrospect about their thought processes in experiments with many trials or to answer general questions, and thus must try to synthesize all the available information after selective recall. In making judgments, subjects have access to an extremely large base of relevant knowledge. Tversky and Kahneman (1973) have demonstrated that subjects only retrieve a few events or pieces of knowledge and use this sample to infer frequencies and probabilities of events. Although the retrieved sample may often be representative and the inferred probability judgment fairly accurate, there are many factors influencing retrievability that do not reflect frequency. Hence, in many situations such cognitive processes will yield incorrect judgments about frequency. Even though all the specific information retrieved is accurate, the inferred probability may be seriously in error. Nisbett and Ross (1980) have given a recent comprehensive discussion of such biasing factors in human judgment.

Particular and General Reports

If the purpose of retrospective probing were to recover memory traces of subjects' processes, the appropriate instruction would be to ask them to recall their specific thought processes during particular trials of the experiment. For at least two different reasons, such a procedure is rarely used. First, after a series of trials, a subject's memory for individual cognitive processes will be poor and lacking in detail. Moreover, there is a tendency for recurrent cognitive processes gradually to become automatic, so that fewer or none of the intermediate states of the processes for the later trials of the experiment are accessible for recall.

Second, many experimenters, because they are interested in general characteristics of the thought processes and not in the episodic details of the individual trials, probe their subjects with questions of the type, "How did you do these tasks?" Such questions implicitly or explicitly request a general rather than specific interpretation of how the subjects were performing the tasks.

There are several different ways in which subjects might arrive at descriptions of their general procedures, as distinct from reports on specific behaviors during individual trials. One possibility is that subjects are aware of the general procedures, or "programs," they are using, use essentially the same programs on all trials, and can recall and report these directly without reference to the specific behavior they produced. Another possibility is that subjects remember some parts of their processes during particular trials, and generalize this information into a general procedure, which they then report. A different possibility is that subjects remember some specific tasks, regenerate–by redoing them–the processes used for these tasks, and use this information to infer the general procedures they may have used. Finally, subjects may draw upon various kinds of prior information, such as general knowledge on how one ought to do these tasks, to generate a verbal report describing a general procedure or strategy. In this case, the verbal reports may not bear any close relation to the actual cognitive processes (Nisbett & Wilson, 1977).

In areas of applied psychological research where verbal questioning has a long tradition, subjects are usually asked about specific events rather than for general information or conclusions. In the critical incident technique proposed by Flanagan (1954), the subjects were always asked to report their memory for specific events. For example:

> ... pilots returning from combat were asked "to think of some occasion during combat flying in which you personally experienced feelings of acute disorientation or strong vertigo." They were then asked to describe what they "saw, heard, or felt that brought on the experience." (Flanagan, 1954, p. 329)

Interpretive probing, unlike the critical incident technique, cannot be relied upon to produce data stemming directly from the subjects' actual sequences of thought processes. The probing procedures encourage or even require subjects to speculate and theorize about their processes, rather than leaving the theory-building part of the enterprise to the experimenter. There is no reason to suppose that the subjects themselves will or can be aware of the limitations of the data they are providing. Moreover, the variety of inference and memory processes that might be involved in producing the reports make them extremely difficult to interpret or to use as behavioral data.

TWO CHALLENGES TO VERBAL REPORTS

It will be useful, in order to get a perspective on the issues, to use the above analysis to examine two published papers that have sometimes been interpreted as providing strong evidence against trusting verbal reports as data from which cognitive processes can be inferred: the first, a paper reporting a study by Verplanck and Oskamp; the second, the review paper on retrospective verbal reports by Nisbett and Wilson. A discussion of these papers will show how the information processing model we have outlined can help us interpret the findings of experiments on verbalization.

Apparent Inadequacies of Concurrent Verbalization

In an often cited study (Verplanck, 1962), Verplanck and Oskamp claimed to have shown that verbalized rules are dissociated from the behavior they were supposed to control. By having subjects verbalize the rules they were following in sorting illustrated cards, the experimenters could reinforce either the verbal rule or the placement of cards (i.e., behavior). To make the contingencies less noticeable, the criterion trials were followed by additional trials with partial reinforcement. When correct placements were reinforced, the subjects were found to place cards correctly in 71.8% of the trials; but they stated a correct or correlated rule in only 48.4% of the trials. When correct statement of the rule was reinforced, the subjects stated a correct or correlated rule on 92.8% of the trials, but placed the cards correctly on only 76.8% of the trials.

In a replication and analysis of this experiment, Dulany and O'Connell (1963) were able to show that the above results could be attributed to two artifacts of the original experiment. First, in the case where correct placement was reinforced, by making a correction for guessing (the subjects had a 50-50 chance of placing the card in the correct pile when they didn't know the rule), we can estimate that subjects *knew* the correct answer in 43.6% of the trials–a percentage very close to the 48.4% in which they stated the correct rule.

Second, with respect to the reinforcement of rules, Dulany and O'Connell found that the rules defined by Verplanck and Oskamp were ambiguous for the card illustrations they employed. In fact, naive subjects who were told these rules explicitly misplaced the cards as frequently as did the subjects in the original experiment.

In a detailed analysis of the rules the subjects verbalized on each trial, Dulany and O'Connell found that on all but 11 of 34,408 trials the subjects put the card where they said they were going to. Hence, Dulany and O'Connell impeached rather thoroughly the evidence put forth by Verplanck and Oskamp for believing that the rules subjects verbalized were inconsistent with their behaviors.

Numerous studies provide positive support for consistency between verbalized rules, concepts, and hypotheses and immediately preceding and succeeding behavior, before subjects receive feedback. When Schwartz (1966) asked subjects their reasons for placing a card as they did, the reasons given were consistent with the placements on all but 2 of 1,962 trials. Even more impressive, Frankel, Levine, and Karpf (1970) obtained retrospective reports from subjects about the basis for their responses to four earlier discrimination-learning problems with 30 nonfeedback trials each, and found that subjects could provide such reports in more than 90% of the sequences of trials.

Apparent Inadequacies of Retrospective Reports

In a recent extensive review of studies permitting evaluation of retrospective verbal reports, Nisbett and Wilson (1977) have reported evidence that appears at first sight to be very damaging to the utility of verbal

reports for inferring information processes. Since their paper has received widespread attention, it is important that we review their findings carefully. The authors summarize their main empirical findings thus (1977, p. 233):

> People often cannot report accurately on the effects of particular stimuli on higher order, inference-based responses. Indeed, sometimes they cannot report on the existence of critical stimuli, sometimes cannot report on the existence of their responses, and sometimes cannot even report that an inferential process of any kind has occurred.

First, we call attention to the frequent use, in their summary, of the qualifiers "often" and "sometimes." Nisbett and Wilson cite a large number of experiments that support their conclusions, but do not investigate in detail the *conditions* under which these conclusions do and do not hold. Moreover, they do not propose a definite model of the cognitive processes as a framework for interpreting the findings they survey. Their theoretical interpretations are entirely informal, resting heavily on an undefined distinction between introspective access to "content" and to "process," or, as they alternatively state it, (1977, p. 255), between access to "private facts" and to "mental processes." Their summary of the kinds of information to which subjects *do* have access is this (1977, p. 255):

> ... we do indeed have direct access to a great storehouse of private knowledge ... The individual knows a host of personal historical facts; he knows the focus of his attention at any given point of time; he knows what his current sensations are and has what almost all psychologists and philosophers would assert to be "knowledge" at least quantitatively superior to that of observers concerning his emotions, evaluations, and plans. Given that the individual does possess a great deal of accurate knowledge ... it becomes less surprising that people would persist in believing that they have, in addition, direct access to their own cognitive processes. The only mystery is why people are so poor at telling the difference between private facts that can be known with near certainty and mental processes to which there may be no access at all.

Nisbett and Wilson also observe that subjects "are often capable of describing intermediate results of a series of mental operations (1977, p. 255)" (i.e., that they hold in STM and can access the symbols that are inputs and outputs to such operations).

We may compare this list of "private facts" and intermediate results that, according to Nisbett and Wilson, *are* accessible to subjects with the kinds of information that our processing model would imply that subjects could report. The individual knows, they say, his focus of attention, his current sensations, his emotions, his evaluations, and his plans. He knows the intermediate results of his mental operations. But these are exactly the kinds of information that, according to our model, would be held in STM and be available for verbal reports.

Unfortunately, the studies reviewed by Nisbett and Wilson provide little data on what information is heeded during the thought processes, and what information is accessible from STM and LTM at the time of the verbal report. Nisbett and Wilson find that the subjects, when *asked questions about their cognitive processes*, frequently do not base their answers on memory for specific events at all, but "theorize" about their processes (1977, p. 233).

> When reporting on the effects of stimuli, people may not interrogate a memory of the cognitive processes that operated on the stimuli; instead, they may base their reports on implicit, a priori theories about the causal connection between stimulus and response.

In reviewing the studies cited by Nisbett and Wilson, we can profitably raise the question of *why* and *when* subjects do not consult their memories of cognitive processes in answering questions about those processes. It is easy to draw the erroneous conclusion that this independence of verbal answers to questions about cognitive processes from the actual course and results of those processes implies a *general* lack of accessible memory for such processes, or even an unawareness of the information while the process was actually going on. But this sweeping conclusion appears not to be justified.

The accuracy of verbal reports depends on the procedures used to elicit them and the relation between the requested information and the actual sequence of heeded information. Invalid reports, like those discussed and obtained by Nisbett and Wilson, may be due to lack of access to thoughts (their claim), inadequate procedures for eliciting verbal reports, or requesting information that could not be provided even if thoughts were accessible. In a subsequent chapter (Chapter 3) we will describe in some detail what information will be heeded and hence reportable. Although some studies cited by Nisbett and Wilson did probe for such information, we will focus here on the deviations between the verbal report procedures used in many of the studies cited by Nisbett and

Wilson and the procedures that, according to our model, would elicit valid retrospective reports of cognitive processes.

First, many of the verbal reports they discuss could be generated without accessing memory of the corresponding cognitive processes. In some of these studies, the questions presented to subjects contain considerable background information from which answers could be generated without consulting their memories. With questions like, "I noticed that you took more shock than average. Why do you suppose you did?" (Nisbett & Wilson, 1977, p. 237) It is not even clear to us, nor probably to the subjects, that memory for the cognitive process *should* be the information source for the answer. If subjects can generate their answers without consulting their memories (Nisbett and Wilson showed that control subjects could do exactly that), they might often prefer this method to retrieving information from memory.

Second, several aspects of the verbal report procedures reviewed by Nisbett and Wilson made the relevant thoughts less *accessible.* In most of the studies reviewed, the time lag between task and probe was sufficiently great to make it unlikely that the relevant information remained in STM. In Chapter 3 we will review the rather extensive literature from general experimental psychology showing that time and intervening thought activity between the cognitive process and its verbal report, as well as incentive to recall memories of the cognitive process, lead to dramatic declines in the accuracy of the verbally reported information. A recent chapter by Genest and Turk (1981) and a paper by Wright and Kriewall (1980) give references showing that such considerations of accessibility are powerful determiners of the accuracy of verbal reports for cognitive processes in tasks like those discussed by Nisbett and Wilson (1977).

A tendency to generate verbal reports without access to memories will be stronger, the less readily available the memory is. When the probe is not a good retrieval cue for the relevant aspects of the memory, the subject must attempt, through conscious processing, to recall sufficient information to give an appropriate answer. Since retrieval from LTM, even if possible, requires considerable time and effort, subjects, unless explicitly instructed to provide a relatively complete recall, may be disinclined to do so, especially if other ways of producing a response are open to them. A recent study by Wright and Rip (1980) provides strong evidence for an increase in accurate self-report when subjects were explicitly motivated to retrieve memory for thoughts in a judgment task.

Finally, in some studies reviewed by Nisbett and Wilson, subjects were asked to report information that cannot be given even with complete access to the thought processes (cf. why-questions regarding causes), and information that is far from a direct recall of memory of the cognitive processes. Our model predicts that information can be recovered by probes only if the same information would be accessed by undirected requests for concurrent or retrospective reports. For many of the studies in the Nisbett-Wilson review, our model would predict failure to obtain from the probes verbal information about particular instances of processes. For example, in between-subject designs, subjects obviously cannot answer from memory of their processes why they behaved differently from subjects in another experimental condition–the processes did not include such a comparison. Hence, this information can be derived, if at all, only by comparing the descriptions of the processes provided by different sets of subjects in the two conditions. In other studies the subjects were asked how they would have reacted if the experimental conditions had been different in a specified respect. Such probing for hypothetical states can never tap subjects' memories for their cognitive processes, since the information was never in memory. In still other studies, subjects were asked, explicitly or implicitly, to summarize or generalize the processes they used, rather than to report concretely the processes used on each trial.

Several articles have been published making similar criticisms of the Nisbett and Wilson (1977) paper, and raising other objections as well. Of particular interest are the papers discussing the problems with verbal reports in between-group designs. (Smith & Miller, 1978). Some recent studies have shown that in corresponding within-group studies, subjects are able to provide veridical verbal reports (White, 1980, Weitz & Wright, 1979; Wright & Rip, 1980).

In sum, we disagree with Nisbett and Wilson's interpretation that subjects simply were not aware of relevant information during the critical experiments. Instead, we claim that better methods for probing for that awareness (concurrent or immediate retrospective reports) would yield considerable insight into the cognitive processes occurring in *most* of the studies discussed by Nisbett and Wilson. On the other hand, we agree with Nisbett and Wilson's analysis of subject's reports in situations where the subjects do not have access to or for other reasons don't rely on memory for the cognitive processes in question. In such situations, Nisbett and Wilson propose that an experimental subject infers the causes of his own behavior by relying on common-sense theories and observable

events–the same process that an observer would use to infer causes of behavior in an observed subject. By using experimental situations, where common-sense theory would lead to the incorrect assessment of causes, Nisbett and Wilson provide convincing evidence for their interpretation by showing that both experimental subjects and observers agree on the incorrect cause of the experimental subjects' behaviors. (For a nice presentation and extension of these arguments see Nisbett and Ross (1980).)

We think that Nisbett and Wilson's paper has been useful in forcing investigators like ourselves to think carefully about the relation of verbal reports to cognitive processes. Many verbal report procedures are justly faulted by their review. However, their results are consistent with our model of concurrent and immediate retrospective reports.

Concluding Remarks

Our examination of two of the most vigorous challenges to the usefulness of verbal reporting leaves intact our belief that such reports– especially concurrent reports, and retrospective reports *of specific cognitive processes*–provide powerful means for gaining information about such processes. The concurrent report reveals the sequence of information heeded by the subject without altering the cognitive process, while other kinds of verbal reports may change these processes. In retrospective reports of specific processes, subjects generally will actually retrieve the trace of the processes. In other forms of retrospective reporting, subjects, instead of recalling this information, may report information that they have inferred or otherwise generated. Hence, in the chapters that follow, we will pay particular attention to the two special forms of reporting–the one concurrent, the other retrospective–that are most likely to yield direct evidence of cognitive processes.

VERBAL REPORTS OF COGNITIVE STATES AND STRUCTURES

Although this book focuses upon cognitive processes, the model and concepts it employs can be extended to the non-cognitive aspects of verbal behaviors. There are several reasons for undertaking such an extension. It will permit us to identify common problems and issues in areas of psychology, like psychophysics, survey design, and measurement of personality traits, that traditionally have had little or no interaction with each other. In these areas, too, as in those we have been discussing, behaviorism has muted explicit examination of the status of verbal responses and reports.

First, we will propose a taxonomy of these other kinds of verbal reports, and will discuss briefly some examples of relevant research. Then we will consider two limited topics for more systematic discussion. The first of these is attitude assessment, the second is the historical development of verbal reporting, with particular emphasis on introspection. All of the verbal reports with which we will be concerned in this section are elicited by probes specifying what information is to be reported. Often, also, a set of alternatives is supplied from which the subject has to select a response.

Predictions from our model about the effects of verbal reporting on thought processes will depend on the circumstances under which the verbalizations are induced. We can classify verbalizations according to the memories that are tapped and according to the verbalization instructions the experimenter gives to the subjects. With respect to the memory source of the reported information, we can distinguish among (a) reports of stimuli that remain constant and available to the subject's senses while the report is being made, (b) reports of information retained in STM, and (c) reports of information from LTM. The next three subsections of this section will be devoted to the special problems that arise for each of these three kinds of reports.

Reporting of Sensory Stimuli

At any given moment, a large amount of external stimulation impinges on any human through the sensory receptors (visual, auditory, etc.), as well as from internal visceral sources. Normally this information is not heeded directly, but recognition processes access existing relevant LTM patterns, which provide higher-level descriptions and are in turn heeded. (In Chapter 3 we will discuss these recognition processes and their relation to attention in some detail.) In many circumstances attention can be directed toward the information in the sensory stores (cf. Kahneman, 1973). We can focus on marks on the page we are reading or listen for unusual faint sounds and so on. Many kinds of verbal reporting procedures rely directly on our ability to process sensory information selectively.