

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER 104-ONR	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) COMPREHENSION AND ANALYSIS OF INFORMATION IN TEXT: IV DECISION AND VERIFICATION PROCESSES		5. TYPE OF REPORT & PERIOD COVERED Technical Report 1/1/81 - 6/30/81 5/1/80 - 5/26/81
		6. PERFORMING ORG. REPORT NUMBER ICS Tech. Report #104
7. AUTHOR(s) Steve Antos; Ely Kozminsky, Ben Gurion University of the Negev, Israel; Lyle E. Bourne, Jr.; Walter Kintsch		8. CONTRACT OR GRANT NUMBER(s) N00014-78-C-0433
9. PERFORMING ORGANIZATION NAME AND ADDRESS Institute of Cognitive Science University of Colorado, Campus Box 345 Boulder, CO 80309		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS NR 157-422
11. CONTROLLING OFFICE NAME AND ADDRESS Personnel and Training Research Programs Office of Naval Research (Code 458) Arlington, VA 22217		12. REPORT DATE May 1981 July
		13. NUMBER OF PAGES 54
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office)		15. SECURITY CLASS. (of this report) Unclassified
		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE
16. DISTRIBUTION STATEMENT (of this Report) Approved for public release; distribution unlimited.		
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)		
18. SUPPLEMENTARY NOTES		
19. KEY WORDS (Continue on reverse side if necessary and identify by block number) Reading, comprehension, memory, decision, fact availability		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) Reading texts for the purpose of making decisions was studied in a laboratory analogue of a complex, natural, information-analytic domain. Subjects, acting as stock brokers were trained in the first two sessions to categorize and evaluate stock report-type information. In addition, they learned to infer information from text-explicit facts. In Session three, subjects read texts and made decisions to Buy or Not Buy based on a conjunctive rule that was either given to the subject before (RB) or after (RA) reading the text.		

In Session 4, subjects read texts and then were presented probes that were to be verified as to either having or not having been presented in the previously read texts. Performance in all tasks was measured in terms of response latency, as well as accuracy. The results demonstrated that a specialized control schema for text comprehension develops that is based on the nature of the decision task. For example, readers in the RA versus the RB condition develop different text analysis techniques that are apparent in the decision, as well as the verification task performance. Most of the results are interpretable in terms of a model which suggests that purposes and goals relevant to the text comprehension process are incorporated into long-term memory (i.e., knowledge structures). The integration of these knowledge structures with current short-term information provides the basis of a task-appropriate text representation.

Abstract

Reading texts for the purpose of making decisions was studied in a laboratory analogue of a complex, natural, information-analytic domain. Subjects, acting as stock brokers were trained in the first two sessions to categorize and evaluate stock report-type information. In addition, they learned to infer information from text-explicit facts. In Session three, subjects read texts and made decisions to Buy or Not Buy based on a conjunctive rule that was either given to the subject before (RB) or after (RA) reading the text. In Session 4, subjects read texts and then were presented probes that were to be verified as to either having or not having been presented in the previously read texts. Performance in all tasks was measured in terms of response latency, as well as accuracy. The results demonstrated that a specialized control schema for text comprehension develops that is based on the nature of the decision task. For example, readers in the RA versus the RB condition develop different text analysis techniques that are apparent in the decision, as well as the verification task performance. Most of the results are interpretable in terms of a model which suggests that purposes and goals relevant to the text comprehension process are incorporated into long-term memory (i.e., knowledge structures). The integration of these knowledge structures with current short-term information provides the basis of a task-appropriate text representation.

Understanding and remembering information for the purpose of making a decision involves a complex set of mental operations. The concept of schema is useful in interpreting these operations. A schema is the coherent mental representation of interrelated knowledge, skills, and intentions that guides the processing and utilization of any pertinent incoming information (Rumelhart and Ortony, 1977). Kozminsky, Bourne, Kintsch (in press) have developed a model for analytic information processing in which schema is a central concept. In their model, schema is said to consist of "...declaritive and procedural information that can act, guide, or modify the information flow in a cognitive system." Of particular concern in the present study is how a schema operates to affect the comprehension of information presented in fairly complex, information rich, technical texts, and controls decisions to be based on those texts.

Our approach in the present study is to isolate and emphasize a single salient characteristic of the schema, viz., its control and specification during reading of what aspects of the text are relevant. Theoretically, this mechanism rejects irrelevant information, allowing a semantic representation of the text to be constructed that is appropriate for subsequent decisions to be made. It is the central goal of this study to determine how several factors common to decision making in technical texts guide the development and use of schemata as control devices that allow a sharpening of the distinction between what is relevant and irrelevant.

The experiments presented here examine the effects of training, task experience, memory and type of decision task on subject's sampling procedures and on the ability of subjects to attend to relevant and disregard irrelevant information in texts. We are interested in the effects of these variables in

complex and (quasi) natural domains. For this purpose, we have developed a stock market environment in which various tasks for subject-analysts can be embedded. The stock market was selected because it appears to be fairly typical of a wide range of interesting, high level decision making domains, including intelligence analysis in the military, data analysis by research scientists, medical diagnosis, and possibly analysis of facts in criminal investigation (Kozminsky, et al., in press).

The first step in establishing a control schema for analysis of stock market reports was to teach the experimental participant to categorize and evaluate single statements that might appear in such a report. Subjects were trained to interpret statements as falling into one of six fundamental information categories regarding any business: Sales, Earnings, Growth, Capitalization, Dividends or General Factors. They also learned to evaluate the degree to which a statement was positive or negative within a given category. Training was intended to build a schema for reading, comprehending and analyzing stock reports. Structurally, the schema has six slots corresponding to six informational categories, each of which has a potential evaluative valence. The slots can be thought of as requests for pertinent evaluative information. Furthermore, the schema operates with two basic processes, categorization and evaluation of any particular input statement. One purpose of reading a specific text, then, is to fill each slot with the appropriate qualitative and evaluative data.

When this schema driven text analysis becomes efficient, the trainee has no need to include in his textual representation any information irrelevant to his purpose. Thus, it is of little consequence to the comprehender what particular product caused Sales or Earnings to explode. It is conceivably

irrelevant for some purposes to know why a company has bad credit as long as it is understood that Capitalization is weak as a result. It is assumed that the processes of categorization and evaluation are fundamental, occurring in a wide range of comprehension tasks which lead to decisions, that they are intimately involved in the development of task-specific schemata and are a basis for distinguishing relevant from irrelevant information. The control schema is applied to arrive at decisions (Kozminsky, et al., in press). In the process, schema knowledge is often used to make inferences about information that might not be presented in a given text. If a report indicates that "Sales are outstanding" but contains no statement about Earnings, the subject might infer, on the basis of what he knows about the relationship between Sales and Earnings that Earnings are also likely to be excellent. Correspondingly, if the report says nothing about Dividends but indicates that the capitaliaztion of the company has improved substantially, the subject might infer that Dividends were probably not high.

In order to examine these inferential processes, participants were instructed and trained to treat certain pairs of categories as perfectly correlated. Two experimental tasks, a completion task and a verification task, were used as a part of the training purposes. The completion task presented the subject with sets of fundamentals and corresponding valences, with the valence of one of the fundamentals missing. The subject was to infer the missing valence, if possible, from existing information about the others. In the verification procedure the subject was to check the set of category-valence pairs for consistency with the predefined correlational structure of categories.

Categories, valences and correlations and the cognitive skills to generate them from text are represented in the general schema which training was intended to develop in all subjects. But particular decision tasks to which the schema can be applied require only a simplified version of the schema. The same general schema then becomes tailored to suit the needs of the specific task. In the general or default schema, all trained categorizations, evaluations, and correlations are executable. As the subject gains experience in any specific task, a sub schema develops that should reduce search time and increase accuracy to the extent permitted by the task demands. This task specific schema, a degenerate from the general schema, will guide search and the extent to which category and evaluative information is processed. Adjustments to that schema made as a function of decision-task specifications produce a more specific control schema. The major control on execution should be in terms of what is considered relevant. This can be realized by a complete cessation of processing when the criterion for decision is reached (e.g., self-terminating search) or reduced level of processing of a particular statement when it is realized that the statement is not relevant to the current needs of the decision rule. If these processing adjustments are governed by a control schema, they will be evident in accuracy and/or RT measures.

In the study by Kozminsky, et al., (in press), subjects acted as stock brokers in a task which required them to discover a certain conjunctive decision rule for predicting performance of a fictitious stock. Subjects read specially designed stock reports and had to learn, by feedback on their decisions, which two categories were relevant to making a correct decision. The task was to identify the relevant categories of information and to acquire

a decision rule based on those categories. In the present study the knowledge of what is relevant is given to subjects. That is, they are given the decision rule that specifies the categories to be considered for a correct response. Whereas Kosminsky, et al. were interested in rule acquisition, the present study is interested in rule usage. The present study was aimed at understanding how the general schema is specialized to control text processing as a function of knowing what specific categorical information is needed to satisfy a decision rule.

One question of major interest is how processing is controlled depending upon whether a person is provided with a decision rule prior to or after reading a text that contains information pertinent to that decision. Subjects make decisions as to whether or not they should buy the stock of a fictitious company. Decisions are to be based upon "quarterly reports" which provide explicit information on three of the six "fundamentals" of the company's recent activities. The decision rule on any give trial joins two different fundamentals which need to be positive for a Buy decision. If the text provides negative information on one or both of the specified fundamentals, the correct decision is "Not Buy". In the Rule Before (RB) condition, participants are given the conjoint rule before reading the text, and in the Rule After (RA) condition the rule is given only after the text has been read.

It is expected that participants in the RB group will incorporate the decision rule into their task specific control schema and deeply process only those parts of the text that are relevant to the rule. Participants in the RA condition will be more likely to process the text completely. Their control schema will be based more generally on their training in earlier sessions (i.e., categorization, evaluation, and correlation) and not further specified

by a particular decision rule. One expected consequence of this difference in processing is that the benefit engendered by a self-terminating search in a "Not Buy" text should be greater for the RB than for the RA condition. The RA group, while it can terminate its search as in the RB condition, has processed other statements in the text and those statements could influence the decision. There is also some evidence to suggest that configural properties of the text will be more influential in the RA condition (Hayes-Roth and Walker, 1979).

Though the above predicts that the RA group will be less able to take full advantage of the processing reduction usually created by self-terminating searches, it is expected that they may enjoy other benefits. There is some evidence to the effect that inferences are made efficiently and probably more accurately on the basis of information that resides in memory than information that is perceptually available (Hayes-Roth and Walker, 1979). That is, when a subject is well-trained in the tasks to be performed, knows the kinds of reports that are going to be read, has a control structure, realizes that certain decisions are going to be required, but has no specific decision rules, then that subject is likely to make all allowable inferences while reading the report. In contrast, when the subject has an explicit decision rule at the outset, the report can be read without necessarily making any inferences at all (when all information required by the decision is explicit in the text). So, when a decision does require inferences, we might expect better performance in decision making in the RA in contrast to the RB condition. This is one of the basic hypotheses to be examined in the present experiment.

As a final task in this experiment, we use a verification procedure to examine the way in which a schema determines how texts are processed and stored in memory. In the verification task, subjects are asked to indicate whether a target statement as true or false in light of information contained in a just-read text. These texts contain, in some cases, both members of a correlated pair. Probes can either be inside or outside the pair. We expect to find inside probes to be verified more accurately and possibly faster than outside probes. This prediction is based on the assumption that correlated categories will augment each other when the related categorical statements are read. It is assumed that verification performance on one of these categories will benefit by the fact that the other one was also in the text. This benefit is only expected to occur, of course, if the possible inferences are made when the text is read. It is also predicted in this verification task that the above mentioned benefit should be greater for subjects who were in the RA condition than for those who were in the RB condition of the decision task. The task-adjusted specialized schema developed during the decision task will be continued to some extent when subjects are placed in a new task. As mentioned above, subjects in the RA condition are likely to process the text as completely as their general schema will allow, making all possible inferences. The RB group learned to more or less make only those inferences dictated by the predefined decision rule. The RA group will be more likely to make the inferences necessary to allow them to benefit from relationships between pairs.

Method

Subjects.

Thirty undergraduate college students participated as subjects in the experiment for four one-hour sessions each. They were paid \$12.00 for their time, plus a \$3.00 bonus ostensibly based on the quality of their performance in the experiment (actually all the subjects received this bonus).

Materials and Apparatus.

Subjects read short texts, variable in length. Mean text length was 47.26 words and the texts range from 24 words to 98 words in length. Each text contained information relevant to three out of six "fundamentals", viz., General Factors, Capitalization, Growth, Sales, Earnings, and Dividends of a fictitious company, EXTEC. The three information categories in each text were selected according to a predetermined sampling.

A pool of 240 sentences, 40 sentences for each of the six information categories, was available for sampling. These sentences were pre-evaluated on a 5-point scale from 1 (negative) to 5 (positive) information about a category (see Kozminsky, Bourne, and Kintsch, (1979)), so that in each category eight sentences represented one of the five possible scale values. Eighteen additional sentences, three sentences per information category, were used in a practice period.

The experiment was controlled by the Sigma 3 computer housed in the computer laboratory for Instruction in Psychological Research in the Department of Psychology, University of Colorado. The instructions and material were presented on CRTs and responses were made on a five button panel interfaced with the computer.

Procedure.

The experiment consisted of four sessions spaced over four consecutive days. Each session lasted about an hour. The first two sessions were training sessions: sentence training and report training. The last two sessions will be referred to as the decision and the value verification tasks. Training.

During the sentence training session, a subject went through the following sequence:

(a) The general nature of the experiment was explained. The subjects were told that they would perform various decision and verification tasks with stock reports. Initially they were trained to read and evaluate the information in such reports in order to maximize performance efficiency in the later tasks. A detailed description of the six information categories was given. (A printed version of this description was in front of the subject for reference throughout the session.)

(b) Eighteen practice sentences were displayed one at a time and the subject learned, with feedback, to assign sentences to their corresponding categories. In a second pass with the same 18 sentences the subject evaluated the information in each sentence on the 5-point scale, again, with feedback.

(c) The nature of decision rules, conjunctions, to be used in the decision task was explained and demonstrated with three tests constructed out of the 18 practice sentences. In addition, a background description of the fictitious company ECKOL was presented and followed by a short comprehension test. While in the first session subjects were primarily trained to identify and evaluate individual sentences, the emphasis in the second report training session shifted to the learning of the report structure to be used in the last two sessions and to the relations among report categories. The sentence

evaluation scale was divided into two zones: a positive zone consisting of the values 4 and 5 on the 5-point scale and a negative zone consisting of the values 1, 2, and 3.

The subjects were told that two categories, Sales and Earnings were positively correlated, so that whenever a sentence value in one category falls into the positive zone the sentence value of the second category will fall into the same, positive zone. If the value of a sentence in one category falls into the negative zone so was the sentence value in the second category. Another pair of categories, Growth and Dividends, was defined to be negatively correlated. A positive zone value in one category entailed a negative zone value in the second category. All other category pairings were uncorrelated.

The correlation concept was carefully explained and illustrated with examples. A "commonsense" explanation for the particular correlation used was given: Sales produce profit and if Earnings are good, one can expect Sales to be good. Similarly, when money goes into Growth, the company has less money to pay Dividends, and when the company pays a large amount of Dividends there is less money to invest in Growth. A typed table describing the correlation among the categories was in front of the subject for reference during this session.

Subjects were given two kinds of experience with the correlation concept. In a completion task, they had to infer missing category values. In a verification task, reports were tested for consistency. Reports were not written out in paragraphs (text form), but rather were given in the form of tables. Each table presented the category labels and their associated values. In the completion task a report consisting of three, four, or five categories was presented in which one of the category values was missing. The subject's

task was to determine the value of the missing category. For example, a presented report was "Earnings - 3, Capitalization - 1, Sale - ?." The subject's task was to infer from his knowledge about the correlation between Earnings and Sales that Sales value was negative. There were three possible responses: the inferred category value was in the positive zone, the negative zone, or could not be inferred. An example of a report that elicited the last response was "Capitalization - 3, Sales - 5, General Factors - 3, Earnings - 4, Growth - ?." Since no information was given on Dividends, the Growth value could not be inferred. Three response buttons from left to right were designated as "-" (negative zone value), "insufficient information," and "+" (positive zone value). The report categories were presented below each other on the screen in a random order except that the last category in the report contained the missing value. Following each response, feedback was given as to whether the subject was correct or wrong and as to what the correct response should be.

In the verification task the subject was presented with a report structure similar to that in the completion task but no category value was missing. The report, again, consisted of three, four, or five categories, and their values. In this task, however, the subject was asked whether the presented report conformed to the correlation structure he was told about. There were three response options: consistent, inconsistent and not inconsistent. An inconsistent report was defined as one that contained at least one correlated category pair, but the category values did not conform to the correlation defined between the categories. e.g., "Growth - 5, General Factors - 4, Dividends - 4." A consistent report contained at least one correlated category pair with values according to the correlation pattern. For

example, "Earnings - 3, General Factors - 2, Capitalization - 3, Sales - 2." The remaining case, not inconsistent reports, contained at least one category from a correlated pair whose counterpart was missing from the report. In the report "Dividends - 4, Capitalization - 3, Earnings - 2, Sales - 1, General Factors - 3," the Growth category was missing. (This case could be thought of as a consistent report with one correlated category missing.) The subject was told that since at least one correlated category was missing there is not enough information to determine whether the report was consistent or inconsistent. The inclusion of this case in the task, forced the subject to test both for consistency and inconsistency of the presented reports. Following each response, feedback was given as to whether the subject was correct or wrong and as to what the correct responses should be.

There were three trial blocks consisting of 18 completion and 18 verification trials each, a total of 36 trials in a block. Of the 18 trials, six contained three categories in a report, six contained four categories and six contained five categories in each sub task. For a six-trial set in the completion task, two trials resulted in a positive response, two in a negative response, and two with an "insufficient information" response. Similarly, in the verification task, there were two consistent, two inconsistent, and two not-inconsistent trials. The number of positively and negatively correlated category pairs were counterbalanced in each 6-trial set. In each 36-trials block, the completion task was presented first, preceded by a message to this effect, and was followed by the 18 trials of the verification task, again, with a preceding message notifying the subject that the verification task was following. Within each 18-trial set, trials were randomly presented and so were the three blocks. The performance emphasis in these tasks was both

accuracy and speed of response. Subjects were told to make their decisions carefully, considering all the available information.

Decision task.

In the third session subjects were asked to read a short stock report and apply a conjunctive decision rule in order to decide whether to buy the stock, not to buy it, or respond that there was not sufficient information in the report to give either one of the first two recommendations. Following each decision, feedback was given on the correctness of the decision. There were 45 such trials. Subjects were assigned to one of two groups, 15 subjects in a group. One group, Rule Before Group, received the decision rule prior to reading the text. The second group, Rule After Group, received the rule after reading the text. Rule inspection, text reading, and decision were self-paced.

The 45 texts were divided into five blocks, nine texts in a block, randomly sequenced for each subject. The five blocks were equivalent in their design. A design configuration for one block is given in Table 1. In the first trial, sentences (and values) from the Earnings (positive), Growth (negative) and General Factors (positive) categories were presented in a text format in a random order. The decision rule was based on two explicit categories in the text - - Earnings and General Factors. The expected correct decision was to buy the stock since the two values of the relevant categories were positive. In Trial 2 Sales (positive), Dividends (positive), and Capitalization (negative) sentences were assembled and presented in a text. The decision rule in this trial was based on Earnings and Dividends. Earnings information was not given in the text so its value had to be inferred from reading the Sales sentence. The other rule category was explicit in the text.

The first three trials in the block example resulted in a buy decision. Trial 1 rule was based on two categories explicit in the text; on Trial 2, one category was explicit in the text and one implicit; on Trial 3, the two rule categories were implicit. Trials 3-6 resulted in a no buy decision with rule category types corresponding to Trials 1-3. The last three trials consisted of texts and decision rules with "insufficient information to make a decision". On Trial 7 the two rule categories were missing in the text and could not be inferred; on trial 8 one category was explicit in the text; and on Trial 9 one category could have been inferred. The other four blocks were equivalent in their design to the one described above, except that the particular categories used and the decision rule varied.

Insert Table 1 about Here

The frequency of categories and their sign values in the texts were controlled over the five blocks. In each block a category appeared in the texts four or five times. About half of the category sign values were positive and half were negative. Each category appeared three times in the decision rule in each block. Texts were assembled from the sentence pool individually for each subject pair (see Materials section), so that one subject from the Rule before Group and one from the Rule after Group received identical texts. For example, a text constructed for Trial 2 in Table 1 consisted of a Sales sentence valued 4 or 5, a Dividends sentence similarly valued, and a Capitalization sentence valued 1, 2, or 3. The order of these sentences in the text was random. Table 2 provides a text sample for Trial 2 in Table 1.

Insert Table 2 about Here

Value Verification Task.

On the fourth session the subject read a short text (self-paced). After reading the text, a category and a value (plus or minus) appeared on the screen and the subject was asked to verify whether the category's value that appeared in the probe matched that in the text. Feedback on accuracy followed each response. In a second task following the feedback, the subject was asked to verify if a category drawn randomly out of the list of the six categories was mentioned in the text or not. No feedback was given after this verification. No data was collected on this second task. It was used merely to insure that the subject processed all the category input and not just the odd-signed category.

There were 48 trials of this kind divided into four blocks of 12 trials each, equivalent in their general design. A design block example is given in Table 3. As an example, on Trial 1 the text consisted of sentences from the Sales, Earnings, and Capitalization categories with negative, negative, and positive values, respectively. The category to be verified was Sales (circled) with a negative sign, so the correct response was "true" to this probe.

Insert Table 3 about here

Two variables were counterbalanced within each block: the expected correct responses and trial type. Three types of text were used: Consistent texts consisted of two correlated categories with signs corresponding to the defined correlations and one uncorrelated category: inconsistent texts in which the sentence signs of a correlated pair were not in accordance with their predefined correlation; and not inconsistent tests in which only one sentence of a correlated pair was allowed in the text. Several other variables were controlled in each block. Each category appeared six times over the 12 trials in each block; a category was probed twice in each block, and there were six category probes with a positive sign and six with a negative sign.

After completing this task, subjects received a questionnaire in which they described their general approach to the task and provided some personal background information.

Results

Training. Response accuracy and latency were measured in both the categorization and evaluation task. There were no significant differences between the Rule Before and the Rule After condition, $t < 1.0$, in either task. The average percent correct out of 18 trials of categorization was 76%, the average latency was 21.0 seconds. The corresponding accuracy and time for the evaluation task were 56% and 9.3 seconds. Thus the two groups of subjects used for RB and RA conditions in later sessions performed about equally in Training Session 1.

The same design was used to analyze the accuracy and latency measures in both the completion and verification analysis: Group x Trial Block x Number of Categories x Response Type. This design yields only two scores per cell. Consequently, for the analysis of correct RTs there were some empty cells. To handle this problem, two different analyses were conducted, one removing subjects with empty data cells and the other collapsing across the number of categories variable. The results were essentially identical in both analyses.

Table 4 shows the main effects and corresponding F-scores. Groups RA and RB did not differ either in percent correct or latency in either the completion or verification task. This is expected since it is not until Session 3 that the treatment manipulation for groups was invoked. The main effect for trial block was also expected, demonstrating significant improvement in performance with practice. There is no significant effect of set size, the number of categories about which information was provided, on accuracy (see Table 4). However, at least for some subjects, increasing the number of categories increased percent correct responses slightly. A positive correlation between set size and percent correct might be due to the fact that large set sizes provide more complete information in a closed task domain. The effect of set size on latency is significant in both the completion and verification task. The more information presented the longer the response latency. Finally, the main effect for response type was significant, but only for latency measures and the verification task. The fact that an effect of response type on RT was found only for the verification task and not for the completion task is understandable in terms of processing differences between the two tasks. First, judging by percent correct and latency measures, the verification task is more difficult than the completion task, $F(1, 20) = 68.4$,

$p < .00001$. The verification task, as mentioned earlier, requires a more complete text analysis than the completion task does.

Insert Table 4 about here

It was predicted that subjects would develop a control schema for these tasks which would allow reduced processing on irrelevant material. Since a greater amount of material was made irrelevant by the completion task, subjects responded faster, on this task. In the completion task there is a set of categories with values and one category without. The subject need only examine which category is with a missing value, search for its correlated category, and compute the missing value, knowing the correlational function. When the correlated category is not present, an exhaustive search of the signed-categories is required. This search model predicts that set size would have its largest effect when the appropriate response is "Insufficient Information". This prediction is substantiated in the data by a significant set size by response type interaction, $F(4, 80) = 4.74$, $p < .002$. It should also be noted, however, that the effect of response type on RT was reduced with practice (Response type by trial block, $F(4, 72) = 3.94$, $p < .01$).

A different search model is required for the verification task. The subject must search the text for information on a correlated category then that category's mate must be found and value relationship between the two must be checked. When there are four or five categories present, two sets of relations might have to be investigated. This explains why responding is more accurate and faster when the appropriate response is "Inconsistent". Only one set of relationships need be found. For "consistent" or "insufficient

Information" response, two sets might have to be dealt with. Further substantiation of this reasoning is found by examining the interaction of set size by response type, $F(4, 80) = 3.19, p < .02$. The largest effect of set size on verification times was when the appropriate response was "Consistent".

The important things to remember about the pattern of results from the training tasks (Session 1 and Session 2) is that it is consistent with task demand models and that Group (RA versus RB) is not involved in any significant effects. These basic results support the notion that a task-adjusted control schema develops with task experience. That schema controls information processing and search procedures. Also, the lack of effects for the RB/RA factor, suggests that any group effects that arise in Session 3 and 4 are most likely attributable to treatment conditions imposed in Session 3.

Decision. The Decision task yields three measures: proportion correct response $[P(c)]$ and two reaction time measures (RT1 and RT2). For the Rule Before condition (RB), RT1 was the time to read the rule and RT2 was the time to read the text plus make a decision. For the Rule After condition (RA), RT1 was the time to read the text and RT2 was the time to read the rule, reflect on the text and make a decision. These RT measures are obviously not directly comparable because a particular measure contains a different amount of reading in each condition. Consequently, RT1 and RT2 were analyzed separately for conditions RB and RA. Except for a tendency for longer RTs for reading texts or rules that contained more implicit information, however, the analyses of RTs showed no significant or interesting differences. The lack of significant effects in RTs is likely due to a high degree of response variability caused in large part by having three as opposed to two possible response categories in conjunction with the complexity of the decision. In any case, most of the

effects of interest obtained in the accuracy scores. This concentration of treatment variability in the p(c) measure was fortunate in that it allowed a better comparison between the RB and RA conditions, since as mentioned above the crucial RT measures for the two groups are not directly comparable.

Table 5 shows the mean percent correct decisions by Response Type and Text Type for each condition. Text type has a different meaning when the response Insufficient Information is required than when either Buy or Not Buy is required. Because of this difference, responses to text containing insufficient information were analyzed separately. Mean proportion correct responses for texts with missing information on both categories specified in the conjoint rule was .74. When information on one category was present, the mean proportion correct responses was .48 if that category was given explicitly and .65 if that information was implicit, $F(2, 40) = 9.9, p < .005$. Apparently, when only one piece of information is provided for a conjoint rule, decision is based on that information if the information is given explicitly as opposed to implicitly in the text. Thus there may be differential weightings of information as a function of how the information is obtained. If one piece of relevant information is text explicit there is a tendency to use it to make a buy/not buy decision rather than to search for the other piece of information required by the conjoint rule. All other main effects and interactions with respect to insufficient information responses were not significant.

Insert Table 5 about here

Response to texts requiring BUY and NOT BUY decisions were analyzed together since the Text Type variable has the same meaning for both Response Types. Mean proportion correct responses for the RB group was .76 and for the RA group .70, $F(1, 20) = 28.2$, $p < .0001$. Accuracy was greater when the text required a NOT BUY decision than when a BUY decision was required. Response type also interacted with condition, $F(1, 20) = 6.1$, $p < .05$. The difference for Response Types was smaller for the RA group than for the RB group.

This interaction reflects the fact that the RA group did not show as great an increase in percent correct in the NOT BUY condition as was found in the RB group. Response Type also interacted with Text Type, $F(2, 40)$, $p < .001$. For texts requiring BUY responses, accuracy increased as a function of the amount of explicit information. For NOT BUY texts, accuracy was largely unaffected by the Text type variable. This possible adjustment in control schema that produces the main effect for Response Type can be understood by the nature of the conjoint decision rule. Both the main effect of response type and the interaction of response type with condition were predicted. It was suggested that a task-adjusted control schema in the decision task would result in a self-terminating search. It was further suggested that because of the RA group having to search through memory, the advantage would be less for them.

Locating one piece of NOT BUY information is enough to support a NOT BUY decision, whereas a decision to BUY needs two pieces of information. The interaction of Response Type by Condition is assumed to be due to the effect of configural text properties that are operational when the information to make decisions is being derived from memory of the text. One possible way this might operate is as follow. The information presented in a sentence is

categorized and memory is searched to determine if that category is relevant to the conjoint decision rule. If the first decision-relevant text category is negative, the appropriate NOT BUY response can be executed without any influence of the second decision-relevant category in the text. In RA processing, the rule categories are read and searched for in memory by the subject. In this condition, even if the first decision-relevant text category found is negative, the valence of the other decision category or possible non-decision categories could affect the decision. Briefly stated, since the RA condition forces processing of the whole text there is a greater chance of positive information inhibiting the NOT BUY response.

The main effect for Text Type was significant, $F(2, 40) = 4.8, p < .01$: decision accuracy was a positive function of amount of explicit information. However, as the above Response type by Text type interaction suggests, the effect is evident only in texts requiring BUY decisions. The increase in difficulty with more implicit categories can be understood in the following way. When the text contains one or more pieces of implicit category information, there is an increased chance of confusion due to explicit and implicit category valences. For example, assume that a rule requires the decision to be based on General Factors and Growth. Suppose that one decision-relevant category in the text was a positive statement about General Factors, and that a second decision-relevant category found in the text was a positive statement about Dividends. In this sort of situation there is an incongruence between explicit and implicit valences. This incongruence occurs about half the time for Text Types containing one implicit category, but always occurs when there are two implicit categories. Thus, when evidence exists for satisfaction of the conjoint rule, differences in valence of

explicit and implicit categories make integration of the true information more difficult. Presumably, such an integration never takes place for a NOT BUY decision, since only one piece of information is necessary.

An analysis was performed to compare proportion correct responses for positive versus negative inferences. This analysis used responses only to texts containing one piece of implicit information. Positive inferences were more accurate than negative inferences, $F(1, 20) = 4.62$, $p < .05$. There was also some indication though not statistically reliable, that inference type was less important for RA than for the RB condition.

Verification. The only effect to reach significance in verification latencies was the main effect for Response Types. "No" responses took longer than "Yes" responses, $F(1, 20) = 9.7$, $p < .01$. This presumably indicates a more exhaustive memory search for "No" responses.

For the initial analysis of accuracy, data for texts that had no correlated categories present were excluded, because the manipulation of probe type has no meaning in this case. The Condition variable that divided subjects into two groups in the Decision task was used as a factor in the analysis of verification data in order to assess the effects of differential pre-verification experience. Mean proportion correct responses for the Condition x Text consistency x Probe type x Response type interaction is shown in Table 6.

Insert Table 6 about here

The Probe type by Condition interaction was significant, $F(1, 20) = 4.4$, $p < .05$. As can be seen in Table 7. Probe type made a difference only in the RA condition, where probing evaluative information inside a correlated pair produced greater accuracy than probing outside.

Insert Table 7-10 about here

As mentioned earlier, this interaction was predicted on the assumption that subjects would carry over the specialized control schemata from the previous decision task in Session 3. In the decision task, RA condition subjects did not know the decision rule until after reading the text. Consequently, they were forced to develop a control schema for text processing which had them making and remembering both implicit (inference) and explicit information. Some evidence for this was found in the RT1 data, which showed that texts with more implicit information took longer to process. If this is the case, subjects in the RA condition should better integrate the information presented on two correlated categories. That is, if positive information was presented for both Sales and Earnings in the same text and the probe was on Sales, memory for the Earnings information represents another path to the required information on Sales. Since the subjects make these inferences it should also be true that they recognize when the value relationship between the correlated categories is incorrect in the text. The counter-assumption for RB condition subjects is that they learned generally only to make those inferences that they needed for collecting evidence on a rule. The verification task provided no rules so that RB subjects would accordingly primarily process the texts for explicit information.

Table 8 shows that across both groups, when the text is consistent, probing inside a correlated pair yields greater accuracy than probing outside, while there is little difference for probe type in inconsistent texts, $F(1, 20) = 7.6, p < .02$. Text type also interacted with Response type, $F(1, 20) = 14.4, p < .002$. Table 9 shows that performance was more accurate on YES responses for inconsistent texts and on NO responses for consistent texts. Across treatment groups, probing inside correlated pairs produced the greatest benefit when the text was consistent with the correlations known by the subject. For example, asking to verify that Dividends were positive would be more accurate if a negative rather than positive statement was made about Capitalization in the text, since the learned correlation between these fundamentals was negative.

The verification task also produced the 3-way interaction shown in Table 10. probe type by Response type by Condition, $F(1, 20) = 5.4, p < .05$. The RB group shows an advantage for probing inside versus outside correlated pairs for YES responses, while the opposite is true for NO responses. For the RA group, performance was always better when the probe was inside correlated pairs and that advantage was slightly greater for NO responses. Previously discussed data from this experiment have substantiated to a large extent that subjects in the RA condition, in carrying over their task-adjusted control schema developed during the Decision task, process the text for implicit information more thoroughly than RB subjects. Consequently, it is not surprising to find all around better performance on inside probes for RA subjects. Those probes, whether requiring a YES or NO response, tap an aspect of the text these subjects took into account while reading the text. The case for the RB subjects is different. They are presumed to have most of the

information stored explicitly so correlated fundamentals are not well integrated for them. It might be further assumed that there is a slight bias for remembering positive as opposed to negative information and/or that searching one's memory of a text to verify a probe is passed primarily on searching for confirming as opposed to infirming evidence. Consequently, when the text is inconsistent and the probe is inside a correlated pair, subjects in the RB group are going to do poorly when a negative response is required.

Assume a text which contains a positive Sales statement and a negative Earnings statement and the probe is positive Earnings. If the above rules are in operation, subjects would search their memory for information that would support positive Earnings. RB subjects should be well practiced at making inferences when they have a specific goal in mind. Consequently, they may find a positive Sales statement sufficient evidence to say Yes, Earnings were positive. This predicts a very low accuracy when NO responses are required for inside probes on inconsistent texts. Table 6 confirms this prediction, which we believe to be the primary basis for the above mentioned 3-way interaction.

Discussion

The experimental tasks reported here were designed to provide data relevant to the processes of information analysis and decision making in texts. It is a major thesis of this study that understanding schema development and application will aid in the interpretation of these processes. We have typified the structure of a schema as a set of requests for specific information. A "request for" something is like knowledge of what to look for, but it is likely that a schema also incorporates "how to" knowledge. That is, what set of operations are necessary to obtain the requested information. The

structural and functional characteristics are probably never static, but are developed and continually evolve through an ongoing means-ends analysis.

During training, subjects learned to analyze, to categorize and to evaluate stock reports. Training leaves schemata that guide the use of categorical and evaluative information in specific tasks. Training provided participants with heuristics applicable over a range of tasks. These heuristics include text analysis only for certain relevant pieces of information identified by task demands. In addition, participants learned to analyze categorized information for its implied as well as its explicit information. Because a person can not represent all aspects of the text after a single reading time the means-ends analysis provides a method for selecting a text representation that has the greatest relevance for current purposes. After a reasonable amount of training it is assumed that the schema developed can be thought of as well formed in permanent or long-term memory.

As new task environments are encountered the structural and functional characteristics of the schema continue to evolve. We examined performance in a task where subjects must use a two category BUY/NOT BUY decision rule. If two (specific) fundamentals are positive, the correct response is "BUY" otherwise "NOT BUY". The control schema operates such that if one of the specified fundamentals is found to be negative a "NOT BUY" response is appropriate without further consideration of other text information. As a consequence, we expect and we found better performance levels for "Not Buy" than for "Buy" texts. We do not claim that this result is critical for our theory, but rather only that such findings are not inconsistent with it. The two relevant fundamentals changed from trial to trial in the decision task. The fact that an advantage for "Not Buy" texts was found in a situation where

the to be considered fundamentals were changing can be viewed as evidence that short-term memory can affect the application of higher-order knowledge to text processing. Heuristics generated by the control schema probably control the level of text processing. For example, a given statement might be processed only to the point of determining its relevancy to the decision rule. If the categorized information in a current text statement is irrelevant then no evaluation analysis is necessary.

This sort of processing heuristic could be present in both the RA and RB groups. However, the difference in task environment between these groups produces differences in schema specialization. The rule in effect provides a sharp criterion for relevancy and several experimental effects substantiated this. Most notably, categorical information irrelevant to the rule affected the RB condition less than the RA condition.

That these effects were the result of a task adjusted control schema is partially explained by analyzing how groups RB and RA processed texts in the verification task. These data show that task relevant processing, developed in one task apparently has somewhat lasting effects. On this basis, we assume that the adjustment in processing is controlled at the knowledge level and not merely at a fast changing level such as at the level of short-term memory.

Briefly, it is assumed that the task specific transfer effects found in the present study tend to support the notion of control schema as a knowledge representation that operates to control processing input by determining information relevancy. There is no sufficient evidence here to allow an assumption that analysis of relevancy is the only or primitive function of a control schema. It does, however, appear that it could be fruitful to think of the schema as a kind of ongoing means-ends analysis that functions to

control the depth of processing of pieces of textual information. From the data presented here it is abundantly clear that some such analysis takes place. It is by no means a new finding that text processing is controlled in part by the readers intentions. Nevertheless, it is useful to integrate this into current conceptualization of executive-type processes in reading for purposes of decision making.

Footnote

This research was supported by a contract from the Office of Naval Research, N00014-78-c-0433, and was conducted in the facilities of the Institute of Cognitive Science, University of Colorado. Requests for copies should be sent to Steve Antos, Department of Psychology, University of Colorado, Boulder, Colorado, 80309.

References

- Hayes-Roth, B. & Walker, C. Configural effects in human memory: The superiority of memory over external information sources as a basis for inference verification. Cognitive Science, 1979, 3, 119-140.
- Kozminsky, E., Kintsch, W., & Bourne, L. E., Jr. Decision Making with texts: Information analysis and schema acquisition. Journal of Experimental Psychology: General, in press.
- Rumelhart, D. E., & Ortony, A. The representation of knowledge in memory. In R. C. Anderson, R. J. Spiro & W. E. Montague (Eds.) Schooling and the acquisition of knowledge. Hillsdale, N.J.: Erlbaum, 1977.

Table 1

Design block example for Experiment 1. The signs "-" and "+" represent categories present in a text and their respective values range. A circle represents a category given in a decision rule.

Trial	Categories in Text and Decision Rule										Rule Category Type			Correct Decision
	S1	Er	Gr	Div	GF	Cap	Explicit	Implicit	Insufficient					
1		⊕	-		⊕		2	0	0	0	0	0	Buy	
2	+	0		⊕		-	1	1	0	0	0	0	Buy	
3	0	+	0	-	-		0	2	0	0	0	0	Buy	
4			⊕		+	⊕	2	0	0	0	0	0	N/Buy	
5	-	0		+		⊕	1	1	0	0	0	0	N/Buy	
6	0	-	+	0		+	0	2	0	0	0	0	N/Buy	
7	0		+	-	0	+	0	0	0	2	2	2	N/Know	
8	-	-	0		⊕		1	0	0	1	0	1	N/Know	
9	+	+	-	0		0	0	1	1	1	1	1	N/Know	

Table 2

Text assembled according to design specification of
Trial 2 in Table 1

Sales of large-scale data processing systems are substantial in dollar terms and are expanding modestly. Increased debt to capital ratio severely reduced Ectex cash position. Directors will meet next month and there is speculation about a stock split of 3 for 1.

Table 3

Design block example for Experiment 2. The signs "-" and "+" represent categories present in the text and their respective value range. A circle represents a category to be verified.

Trial	S1	Er	Gr	Div	GF	Cap	Verified Category Sign	Correct Response	Trial Type
1	0	-				+	-	Yes	Consistent
2			+	0	+		+	No	Consistent
3	-	0				+	+	Yes	Inconsistent
4			0	+	-		-	No	Inconsistent
5			-	+	0		-	Yes	Consistent
6	+	+				0	+	No	Consistent
7		0	-	-			+	Yes	Inconsistent
8	-	-			0		-	No	Inconsistent
9	+		-			0	-	Yes	Not Inconsistent
10				0	+	+	+	No	Not Inconsistent
11		-	0		-		+	Yes	Not Inconsistent
12	0			+		-	-	No	Not Inconsistent

Table 4

	Completion		Verification	
GROUP				
RULE BEFORE	.980 (3.85)	F(1,20)=3.7	.875 (6.89)	F(1,20)=2.7
RULE AFTER	.953 (4.42)	[F(1,20)=1.2]	.907(6.48)	[F<1]
FIRST TRIAL	.946 (5.10)	F(2,40)=5.1*	.871 (7.74)	F(2,40)=5.1*
SECOND BLK	.971 (3.73)	[F(2,36)=27.8]**	.939 (6.13)	[F(2,22)=20.4]**
THIRD	.982 (3.08)		.909 (5.34)	
THREE	.968 (3.68)	F(2,40)=1.47	.919 (5.77)	F<1
FOUR	.957 (4.09)	[F(2,40)=19.0]**	.912 (6.30)	[F(2,40)=18.6]**
FIVE	.974 (4.67)		.889 (7.37)	
RESPONSE TYPE				
NEGATIVE (Inconsistent)	.954 (4.12)	F(2,40)=2.8	.934 (5.92)	F(2,40)=2.6
INSUFFICIENT INFO (Not Inconsistent)	.985 (4.11)	[F<1]	.907 (6.90)	[F(2,40)=5.5]*
POSITIVE (Consistent)	.960 (4.21)		.879 (6.62)	

*.01

**.0001

Note: RT averages and F-values are in square brackets. The level labels for response types enclosed in parentheses are for the verification task. Those not in parentheses are for the completion task.

Table 5

Percent Correct Response for Various Conditions in the Decision Task

Treatment Condition	Number of		Rule Before		Rule After		
	Expl*	Impl*	Buy	Not Buy	Buy	Not Buy	Insf Info**
1	0	2	.46	.89	.56	.71	.69
2	1	1	.67	.84	.64	.78	.49
3	2	0	.87	.82	.80	.71	.64

* This information is relevant for BUY/NOT BUY decisions only

** For Insufficient Information

Treat Cond. #1 = 2 missing/

#2 = 1 missing/1 explicit

#3 = 1 missing/1 implicit

Table 6

Percent Correct Response for Various Conditions in the Verification Task

RESPONSE TYPE	PROBE TYPE	RULE BEFORE			RULE AFTER		
		CON	INCON	NOT INCON	CON	INCON	NOT INCON
YES	Inside	.82	.86	—	.80	.77	—
	Outside	.61	.93	—	.66	.80	—
	No	—	—	—	—	—	.76
	Correlated	—	—	.82	—	—	—
NO	Inside	.84	.59	—	.89	.80	—
	Outside	.84	.79	—	.77	.66	—
	No	—	—	—	—	—	.83
	Correlated	—	—	.80	—	—	—

Table 7

Percent Correct Response in Conditions RA and RB for Probes
Within and Outside a Correlated Pair of Categories

		Probe Type	
		IN	OUT
Condition	RB	.78	.80
	RA	.81	.72

Table 8

Percent Correct Response for Consistent and Inconsistent Texts
and Within and Outside a Correlated Pair

		Probe Type	
		IN	Out
Text	CON	.84	.72
	ICON	.76	.80

Table 9
 Percent Correct Positive and Negative Responses
 in Consistent and Inconsistent Texts

		Response	
		YES	NO
Text	CON	.72	.84
	ICON	.84	.71

Table 10
 Percent Correct Response for Different Probes,
 Conditions and Responses

	Rule Before			Rule After		
	YES	NO		YES	NO	
Probe Type	INSIDE	.84	.72	INSIDE	.78	.84
	OUTSIDE	.77	.82	OUTSIDE	.73	.72