Chapter 1

Metacognition in Strategy Selection

Giving Consciousness Too Much Credit

Melanie Cary and Lynne M. Reder
Carnegie Mellon University, USA

Key themes: Cont / Defn / Prob / Proc / Theor

Key words: Metacognition / Strategy selection / Cognitive processes / Adaptivity / Implicit learning

Abstract: Many researchers believe that metacognitive processes regulate strategy selection. Another common assumption is that metacognitive processes, such as strategy selection, entail conscious processing or decision making. In this chapter, we examine whether conscious awareness is a critical aspect of strategy selection. We review evidence that first establishes that strategy selection varies both across and within individuals in response to dynamic features of the environment. Then, we present evidence that strategy adaptation can occur without (a) conscious consideration of different strategies or (b) conscious awareness of factors influencing one's strategy use. Specifically, shifts in strategy use occurred when people seemed to be unaware (a) that there were shifts in their strategy use or (b) that there were changes in the characteristics of the environment that, nonetheless, affected their strategy use.

Conscious awareness is generally believed to be a necessary condition of metacognition (e.g., Paris, this volume). Some researchers clearly express this belief, as evidenced by the recent statement that "metacognition is an essentially conscious activity" (Darling, Sala, Gray, & Trivelli, 1998, p. 89). Many others tacitly hold this assumption, as noted by Tulving (1994). However, a few researchers have recently proposed an alternative view of the role of consciousness in metacognition by suggesting that there are metacognitive processes that do not depend on conscious awareness (Koriat, 2000; Sun, 2000). From this perspective, metacognition involves a combination of explicit and implicit processes.

Additionally, metacognitive processes have frequently been proposed as the mechanisms that regulate strategy selection (e.g., Dunlosky & Hertzog, 1998; Nelson & Narens, 1990; Roberts & Erdos, 1993). In this chapter, we examine whether conscious awareness is a critical aspect of strategy selection. In doing so, we establish that strategy selection varies across and within individuals in response to dynamic features of the environment. We present evidence that this strategy adaptation can occur without awareness of the environmental changes that affect adaptation and without awareness that shifts in strategy have occurred. Therefore, we contend that if metacognition is dependent on conscious awareness, then not all strategy selection involves metacognitive processes.

1. METACOGNITION AND STRATEGY SELECTION

Consistent with the view that metacognitive processes are involved in strategy selection, several researchers have suggested that two principle components of metacognition are monitoring and control of cognitive processes (e.g., Metcalfe, 1996; Nelson, 1996). This conception of metacognition underlies the edited book "Metacognition and cognitive neuropsychology: Monitoring and control processes" (Mazzoni & Nelson, 1998). Monitoring involves assessing information about one's knowledge and performance. Monitoring processes do not inherently require conscious awareness, but the term "metacognitive monitoring" typically refers to processes that do. Control involves self-regulative processes that direct and modify one's behavior, such as processes that govern the selection of strategies for accomplishing tasks.

Monitoring and control of cognitive processes are central to everyday functioning. Consider the processes involved in selecting which procedure to use to calculate a tip. A person can calculate the tip mentally, use paper and pencil, or use a calculator. Further, the person must initiate the processes responsible for each step in obtaining the tip (e.g., round up bill total to the nearest dollar, locate calculator). Understanding how it is that people self-regulate and how these control processes are affected by monitoring is important for understanding human behavior and cognition. Do people shift among available strategies or processes to perform a task? Do people evaluate how they are doing while performing a task or a series of tasks? Do people measure what they know in order to select the most appropriate strategy?

Monitoring and control processes are almost certainly interdependent. Some metacognitive researchers have proposed that there is a causal link between metacognitive monitoring and metacognitive control (e.g., Barnes, Nelson, Dunlosky, Mazzoni, & Narens, 1999). Generally, monitoring processes must exist, because strategy use has been shown to be influenced by the prior success of each strategy (e.g., Reder & Schunn, 1999). The information gleaned from monitoring may be used later to facilitate selection of the best strategy when the person re-encounters a problem of the same type. Consider children's selection of a strategy for solving simple addition problems. Siegler (1987) showed that young elementary school children use a variety of addition strategies and do so in a way that is adaptive to different types of problems. Children tend to use the strategy that produces the most beneficial combination of speed and accuracy for a particular problem. Specifically, children choose faster strategies (e.g., retrieval) when these strategies produce a correct answer. They choose slower strategies (e.g., counting) when these strategies yield a correct answer and the faster ones produce errors. Presumably, cognitive processes monitor the speed and accuracy of each addition strategy by problem type, and this information influences which strategy is used on a particular problem. The question arises as to whether a child's selection of a strategy results from deliberate consideration of the choices and conscious awareness of their prior success rates or from more autonomous, implicit processes. Is conscious control a general requirement of strategy selection? Must an individual be aware of the prior history of success with each strategy?

1.1 Overview and organization of the chapter

We argue that a great deal of strategy selection happens without conscious deliberation or awareness of factors influencing one's choice. The term "strategy selection" may seem to suggest that the process involves a type of deliberate choice, but we believe this process often lacks conscious deliberation. Therefore, we also contend that if *metacognition* is understood as a property of mind that requires conscious awareness, then much of cognitive monitoring and control occurs without metacognitive involvement.

It is important to note that in this chapter we do not distinguish between strategies and procedures. Readers may believe that it would be more appropriate to refer to some of the strategies that we discuss as procedures. Both of these terms refer to means that people may use for accomplishing a task. We are interested in the nature of the mechanisms that result in an individual using one of multiple means for accomplishing a task, regardless of whether the alternatives are procedures or strategies. Furthermore, we believe that the presented effects and conclusions hold for both strategies and procedures.

The rest of this chapter is organized in the following manner: First, we establish the need to posit strategy selection in performance by illustrating that there is strategy variability in people's performance of multiple tasks, not only between but also within individuals. Second, we discuss factors that affect strategy selection, focusing on two types of factors that were proposed by Reder (1987, 1988). In examining these factors we demonstrate that strategy variability can be caused by sensitivity to features in the task and the task environment. Where possible, we report participants' levels of conscious awareness of changes in the task features and their shifts in strategy. Lastly, we consider more generally the necessity of conscious awareness for strategy adaptation. We will conclude that strategy selection can occur implicitly, without conscious consideration of alternative strategies.

2. STRATEGY VARIABILITY

Until relatively recently, the accepted view of performance assumed that the cognitive scientist was to identify *the* procedure used to perform a task. Developmental psychologists assumed sequential or stage-like use of alternative strategies, with more sophisticated strategies being adopted later. More recently, people have provided evidence for the notion of strategy variation within a person performing a task (e.g., Erickson & Kruschke, 1998; Lemaire & Reder, 1999; Lovett & Anderson, 1996; Reder, 1982, 1987, 1988; Siegler, 1988). Consider strategy selection in question answering. In the 1970s and 1980s, the dominant view was that people *first* search memory for the answer when a question is posed. If, and *only* if, that search fails, then people attempt an alternative strategy. This type of view was apparent in several theoretical perspectives, such as SOAR (Laird, Newell & Rosenbloom, 1987) and the Distribution of Associations model (Siegler & Shrager, 1984).

Early evidence that people vary in question-answering strategies was provided in Reder (1982). Participants in that study read brief stories and then were asked to make judgments about each of a series of statements that, based on the stories, were highly plausible, moderately plausible, or implausible. Some participants were asked to discriminate previously read assertions from statements that were not part of the story. These participants were only tested with highly plausible and moderately plausible statements. Other participants were asked to judge the plausibility of each test statement. In addition to the judgment task manipulation, Reder also manipulated the delay between when participants read a story and when they were tested on it: Testing was immediate, 20 min later, or 2 days later.

The results of that study indicated that people can use either a direct retrieval strategy or a plausibility strategy and that people do not always try direct retrieval first. When judgments are made immediately after reading a story, the statements in the story should be relatively accessible from memory; however, after a 2-day delay the information should be less accessible. Participants who were tested immediately were much more likely to first try direct retrieval than participants tested with a 20-min or a 2-day delay. Importantly, this was true regardless of whether participants were asked to judge plausibility or asked to make recognition judgments. Participants who were tested after a 2-day delay had a tendency to try a plausibility strategy first, using retrieval as more of a backup strategy. In sum, Reder (1982) demonstrated that the two judgment strategies had shifting propensities for individuals to try a particular one first.

There is a variety of other evidence for strategy alternation within individuals, some of which will be discussed in more detail below. Reder (1987, 1988) provided evidence that for each test statement there is a strategy selection phase in which individuals making story fact verifications select among plausibility and direct retrieval. Strategy variation has also been documented in an air traffic control task (Reder & Schunn, 1999; Schunn & Reder, 1998), as well as in the problem solving and arithmetic domains. Lovett and Anderson (1996) had participants perform a problem-solving task in which there are two alternative strategies available for the first step in attempting a solution. They found that individuals solving a series of these problems varied their use of the two alternative strategies, rather than each person only using one strategy across several problems. Lemaire and Reder (1999) examined strategy selection in an arithmetic verification task. Again, individuals used a variety of strategies when performing multiple verifications. Because the selection of a strategy for each verification problem was influenced by features of the problem, an individual's strategy use varied across problems. As mentioned earlier, children solving simple arithmetic problems sometimes achieve the answer by retrieval and other times by using one of a variety of different strategies, such as adding by counting up from the larger number (Siegler, 1987, 1988). Variation in strategies for simple arithmetic is not limited to children. Adults performing a running arithmetic task that required multiple additions used retrieval on some steps and counting on others (Cary & Carlson, 1995). Furthermore, individuals appeared to fluently switch between addition strategies without deliberately choosing one strategy over another.

3. FACTORS AFFECTING STRATEGY SELECTION

Given that we have established that people select among multiple strategies for task performance, and they seem to vary in their preferences within the same task, can we specify the factors that influence the strategy selection process? For example, in question answering what determines whether a person searches memory for an answer or infers an answer? Reder (1987, 1988) emphasized two types of factors that affect strategy selection: intrinsic and extrinsic. Intrinsic factors involve people's familiarity with features of the task, problem or question, such as familiarity with terms in a problem. Extrinsic factors involve features of the general context in which a task is performed, such as task instructions and prior history of success with a strategy.

3.1 Intrinsic factors

Many of the experiments described in Reder's (1982) work could be explained by postulating that a person shifts strategy preference before reading the question. Experiments that involve a change in base rates of success, specific advice about the best procedure for the next question, or a delay between study and test all allow participants to decide *a priori* which strategy is likely to prove more efficient. In real life people do not necessarily know the age of the relevant memories when they are queried. Therefore, it might be reasonable to assume that the decision to use plausible reasoning rather than searching memory is based on attributes of the question rather than the tacit knowledge "I have come back 2 days later and thus the information must be old."

A study reported by Reder (1988) provides evidence that people's strategy selection can be influenced by intrinsic features of a task; in this case features of a test statement. The study consisted of two experimental sessions that occurred 2 days apart. Participants read 5 stories in the first session and 5 stories in the second session. After each story in the second session, participants were tested on the immediately preceding story and one of the stories read 2 days earlier. At the beginning of each test, participants were informed as to which story would be tested from 2 days prior. Because the test statements were presented randomly, rather than blocked by target story, participants did not know which of the two stories a statement was related to until after they read the statement. Thus, a critical difference between this study and the Reder (1982) study discussed previously is whether or not participants would know the age of the relevant material prior to reading the test statement. As in the earlier study, half of the participants made recognition judgments and half made plausibility judgments.

The results of that study indicated that participants in both the recognition and plausibility conditions used both direct retrieval and plausibility strategies to make their judgments. Importantly, the selection of a strategy was influenced by the age of the relevant story. When test statements were related to the immediately preceding story, the data indicated that participants tended to use the direct retrieval strategy. However, when test statements were related to a story from 2 days earlier, participants tended to use the plausibility strategy. The data also indicated that participants did not always try retrieval before using the plausibility strategy. Hence, participants were able to adjust or make their strategy choice after seeing the test statement, and the strategy selected was influenced by an intrinsic feature of the probe. These results suggest that when a problem is presented there is an initial evaluation phase that influences strategy use.

When a person initially evaluates a question or problem, how do intrinsic factors influence which solution strategy is tried first? In the case of question answering, by what means does age or familiarity of the material lead a person to decide to use a plausibility strategy first rather than search memory first? We argue that one of the criteria used in this decision is a quick *feeling of knowing* or familiarity that comes from features of the question or test statement. A strong feeling of knowing is likely to lead someone to try retrieval first, while a weak feeling of knowing is likely to result in the use of some other strategy. Furthermore, this feeling of knowing is not based on searching for the answer (Reder, 1987, 1988; Reder & Ritter, 1992). There is evidence that people can know whether they will be able to answer a question *before* they can find the answer in memory. The first evidence for these claims came from studies that used a "game show" paradigm.

In non-experimental settings, most people have observed that television game show contestants appear able to indicate whether or not they know an answer before they have retrieved it from memory, often before the entire question is read to them. Conceivably all people can do this on a more regular basis. Reder (1987) tested this idea by using world knowledge questions in a paradigm intended to be treated like a game show. Participants were asked questions like "Which ship carried the Pilgrims to America in 1620?" and "What was the name of the clown on the *Howdy Doody* television show?" Half of the participants were asked to estimate whether or not they thought they would be able to generate an answer for each question. When they answered "yes" they were then asked to come up with an answer. The other half of the participants were simply asked to answer each question. All participants were encouraged to respond as rapidly as possible. Participants in the answer condition began articulating the answer more slowly than participants in the estimate condition indicated whether or not they thought they could answer the question. Participants in the answer condition were also slower to respond "don't know." This effect was not due to a speed-accuracy tradeoff, because participants in the answer condition correctly answered 74% of the questions they attempted, and participants in the estimate condition correctly answered 88% of the questions they attempted. A control experiment that required participants in both conditions to press a key also found that participants in the answer condition responded more slowly than those in the estimate condition. These data are consistent with the idea that there is a feeling-of-knowing mechanism that allows people to evaluate whether or not they are likely to know the answer to a question before they can actually answer the question.

Conceivably the ability to rapidly evaluate one's knowledge about a question could derive not from a rapid feeling of knowing, but rather from an early stage of retrieval (see Nhouyvanisvong & Reder, 1998; Miner & Reder, 1994 for reviews). Reder and colleagues (Reder, 1987; Reder & Ritter, 1992; Reder & Schunn, 1996; Schunn, Reder, Nhouyvanisvong, Richards, & Stroffolino, 1997) conducted experiments to try to rule out this alternative. They manipulated feeling of knowing by manipulating participants' familiarity with intrinsic features of the questions or problems. Generally when a target item has been seen recently or frequently, familiarity with the item should be high and the item should be relatively accessible from memory. Hence, items with high familiarity should produce a high feeling of knowing.

Reder (1987) manipulated familiarity with test questions in the game show paradigm by priming terms in the questions. For one-third of the game-show questions two terms in each question were previously rated by the participants for co-occurrence (e.g., "How often do golf and par appear together?"). Question difficulty was also manipulated by dividing questions into three levels of difficulty based on norms developed by Nelson and Narens (1980). As in the preceding experiment, half of the participants estimated whether they knew the answer and the others simply answered the question if they could.

As predicted, questions that contained primed terms gave people the impression that they knew the answer to questions that they could not answer. Easier questions should already afford a high feeling of knowing and, thus, priming cannot and did not raise the tendency for participants to think that they knew the answer. However, for harder questions, participants in the estimate condition judged that they could answer more primed questions than unprimed questions. Their ability to correctly answer those questions did not increase with their first impression that they knew the answer. In contrast, for participants in the answer condition there was no effect of priming on proportion of questions attempted. Rather, the effect of priming appeared in the time to respond "don't know" to questions that were not attempted. Participants in the answer condition were slower to respond "don't know" for primed questions than for unprimed questions, suggesting that they were searching longer for an answer to primed questions. These results support the ideas that feeling of knowing is influenced by recent exposure to features of the target item and that this feeling affects whether and how long people search for the answer. Research by Metcalfe, Schwartz, and Joaquim, (1993; Schwartz & Metcalfe, 1992) also emphasizes the role of familiarity with question features. They found that priming of cues in a paired-associate study increased feeling of knowing without affecting memory performance.

Feeling of knowing mechanisms can also be used for initial strategy selection in arithmetic verification problems (Reder & Ritter, 1992). Participants solved "sharp" problems and either addition (Experiment 1) or multiplication (Experiment 2) problems, such as 14#17, 37+15, and 19*13. The "sharp" operator (#) was designed to be fairly equivalent to multiplication in computational difficulty. Participants were presented with a large number of problems, many of which were repeated up to 20 times. When each problem was presented, participants had approximately three-fourths of a second to select between calculating or retrieving the answer. Participants received a 10 times greater reward for selecting retrieval if they could answer correctly within 2 sec. Occasionally participants were presented with operator switch problems in which a practiced problem had its operator switched to a different operator. For example, 19*13 may have been presented several times and then for the first time 19#13 was presented. There are a few important aspects of the results. First, frequency of exposure affected participants' tendency to believe that questions could be answered by retrieval. With each re-presentation of a problem, participants were more likely to select retrieval for solving that problem. They also became faster at correctly answering the problem. Critically, as participants acquired experience with a problem they were more likely to select retrieval when an operator switch problem was presented. In fact, regression analyses indicated that it was experience with the operand pairs, not the entire problem, that predicted tendency to select retrieval. These results suggest that familiarity with terms of the problem, rather than the ability to retrieve the answer, drives feeling of knowing.

Additional evidence indicates that strategy selection was not due to information from an early stage of retrieval (Reder & Ritter, 1992; Reder & Schunn, 1996; Schunn et al., 1997). Reder and colleagues examined the possibility that operator switch problems gave a high feeling of knowing because the wrong answer was being retrieved in from memory. Some evidence against this alternative account comes from Reder and Ritter's study. Their participants were much less likely to choose retrieval for operand reversal problems (e.g., 23x12 vs. 12x23) than for previously encountered problems or for operator switch problems. In other words, their participants failed to indicate a feeling of knowing for problems to which they did know the answer when the problems looked different, while at the same time they quickly responded retrieve, indicating that they felt that they knew the answer, to problems that looked familiar even though they had never been seen before.

More recent work experimentally de-coupled familiarity with the arithmetic operands from familiarity with the answer (Schunn, et al., 1997). This was accomplished by blocking participants from calculating the answers 5 times out of 7 presentations for a set of special problems. For these problems participants made the initial retrieve-compute decision, but instead of providing an answer on most of the trials the screen was cleared, and the participant was instructed to continue on to the next problem. To ensure that participants could not learn that the special problems were never answered, participants were occasionally required to answer them. These infrequently answered problems were one-fourth of all problems in the experiment. In other respects, the experiment resembled the Reder and Ritter (1992) study. As before, participants were presented with a series of multiplication and sharp problems with repetition of individual problems. On each trial, participants made a rapid decision as to whether they would retrieve or calculate the answer. By frequently blocking participants from solving the special problems, familiarity with the problems and learning of the answers were different than for regular problems.

If feeling of knowing is due to an early stage of retrieval, then participants should rarely select retrieval for the infrequently answered problems, because the answer is not associated with the problem on most trials. However, if this feeling is based on familiarity with features of the problem, then participants should select retrieval for the infrequently answered problems as a function of the amount of exposure to the problem, because problem familiarity is increased with each exposure. The predicted pattern of strategy selection for the normal problems (i.e., the ones answered each time they are presented) is the same under each hypothesis: Participants should select retrieval for the normal problems as a function of the amount of exposure to the problem. As predicted by the familiarity hypothesis, for both infrequently answered problems and normal problems the probability of selecting retrieve was a function of exposure, and the same function appeared to underlie strategy selection in both problem conditions. Thus, the tendency to select retrieval was affected by exposure to the problem itself and was not dependent upon exposure to the answer. This result provides strong evidence for the role of the intrinsic factor of problem familiarity in initial strategy selection.

There are other task domains in which intrinsic features have been shown to affect strategy choice. Problem schemata for word algebra problems have been cued by superficial problem features (Hinsley, Hayes, & Simon, 1977). Strategy selection in problem solving has been influenced by how close in absolute distance each strategy will get one to the goal state (e.g., Lovett & Anderson, 1996). Analogical remindings are also influenced by superficial similarity between the current problem and the remembered analogy (Ross, 1984, 1989).

3.2 Extrinsic factors

There are a variety of extrinsic factors that also affect strategy selection, including task instructions, advice about what strategy will work, the availability of working memory support, and prior history of success with a strategy. For example, Cary and Carlson (1999) demonstrated that the strategy that people use to solve a complex problem is influenced both by the availability of working memory aid and the availability of a worked example problem. Participants could use one of multiple strategies to solve each of several income calculation problems. When an example problem was provided, participants were more likely to select the illustrated solution strategy than when no example was provided. Participants provided with working memory aid, via the availability of paper and pen, tended to settle on using a problem-solving strategy that corresponded with the conceptual structure of the task, whereas participants without this memory aid tended to settle on using a different strategy, one that minimized demands on working memory.

A strategy's past history of success appears to be a primary extrinsic factor that affects strategy selection. The influence of this factor has been shown to affect strategy choice in several domains, including runway selection in an air traffic control task (Reder & Schunn, 1999), story question answering (Reder, 1987), equation verification (Lemaire & Reder, 1999), and simple problem solving (Lovett & Anderson, 1996). In general, people tend to be sensitive to base rates of success and select a strategy based in part on what has been successful in the past.

Reder and Schunn (1999) studied strategy adaptivity to changing base rates using an air traffic control task. Participants had to land planes selecting between a short and long runway. There are various rules for the task including rules that govern when a short runway can be used for different types of planes. A long runway can always be used for all planes, but 747s always require a long runway. When most of the planes that must be landed are 747s, participants should try to avoid using the long runway for smaller planes. When there are few 747s, use of the long runway for short planes does not matter, and it is more efficient to land all planes on the long runway. To investigate whether people adapt in this way, Reder and Schunn varied the proportion of 747s in different blocks of the experiment. For example, for some participants the proportion of 747s was 25% in the first block of trials, 5% in the second block, and 50% in the third block. To assess strategy adaptivity, they looked at how frequently participants selected the short runway when both runways were open and could be used. They found that participants generally adapted in response to the proportion of 747s and did so in the expected directions: When the proportion of 747s increased participants increased their use of the short runway, and when this proportion decreased participants decreased their use of the short runway. Interestingly, strategy adaptivity varied across individuals with some participants being more adaptive than others. Furthermore, inductive reasoning and working memory capacity were positively correlated with individual differences in adaptivity (see Reder & Schunn, 1999; Schunn & Reder 1998 for more information).

In one of the experiments described by Reder (1987), participants judged the plausibility of statements based on the story they had just read. Some of the plausible test statements had been presented earlier as part of the story and some of the implausible test statements were identical to statements in the story except for one word whose opposite meaning had been substituted. Thus, participants could use either a retrieval strategy or a plausibility strategy to make judgments. The proportion of test statements that could be judged based on an explicitly presented statement in the story varied both across participants and across the experiment. For the first 6 of 10 stories, the successfulness of the retrieval strategy differed for the two groups of participants. For participants in the direct retrieval bias condition, 80% of the plausible statements had been presented in the stories and 80% of the implausible statements were contradictions of statements that had been presented. In contrast, for participants in the plausibility bias group only 20% of the plausible statements had been presented and 20% of the implausible statements were contradictions of presented statements. For the last 4 stories in the experiment, the percent of statements that could be verified using direct retrieval shifted to 50% for both groups.

Direct retrieval should be a faster process than judging plausibility at the short delay used in this study, and direct retrieval can only occur when the test statement or its contradiction had been presented in a story. Hence, Reder examined the difference in response time for stated probes and not-stated probes. During the bias manipulation (i.e., stories 1-6), this difference between stated and not-stated probes was much greater for participants in the direct retrieval bias condition than for participants in the plausibility bias condition. These data indicate that participants in the retrieval bias condition used the retrieval strategy, even though the task required judging plausibility, whereas participants in the plausibility bias condition were much less inclined to use retrieval. When the proportion of previously presented test statements changed to 50%, the response time difference between stated and not-stated probes was equivalent for the two groups of participants. That is, from the first to second part of the study, the response time benefit for previously presented statements decreased for participants in the direct retrieval bias condition and increased for participants in the plausibility bias condition.

Furthermore, when the data are analyzed with regard to the difference in response time between moderately and highly plausible statements, the data indicate a complementary shift in the tendency to use the plausibility strategy. During the bias manipulation, the response time benefit for highly plausible statements was larger for participants in the plausibility bias condition. When the bias was removed, this benefit decreased for participants in the plausibility bias condition and increased for participants in the retrieval bias condition. Hence, the data strongly suggest that as the retrieval strategy became less successful for participants in the direct retrieval bias group they became less likely to use it, and as the retrieval strategy became more successful for participants in the plausibility bias group they became more likely to use it. These results are consistent with the idea that people are sensitive to base rates of success with a strategy and this sensitivity can produce changes in the likelihood of selecting a particular strategy.

Participants in that study were questioned at the end of the experiment regarding their level of awareness. They were not consciously aware of either the strategies that they were using or the different base rates of success for using retrieval. Regardless of bias condition, the participants thought that they had used direct retrieval to judge plausibility. They also did not differ in their guesses of base-rate as a function of condition.

Lemaire and Reder (1999) found evidence that strategy selection in arithmetic verification can be sensitive to a strategy's base rate of success without participants conscious awareness of these base rates. Their participants were presented with a series of correct and incorrect multiplication problems and asked to indicate whether each one was true (e.g., 6 x 32 = 192) or false (e.g., 8 x 7 = 58). There were two kinds of false problems. Parity match problems had a false answer with an odd-even status the same as that of the correct answer (e.g., 8 x 7 = 58), whereas parity mismatch problems had a false answer with an odd-even status different from the correct answer (e.g., 8 x 7 = 57). Participants were faster at rejecting parity mismatch problems than parity match problems, indicating that people can use violations of the parity rule to rapidly reject false problems. In other words, participants sometimes used a parity-check strategy to reject problems. Lemaire and Reder manipulated the proportion of false problems violating the parity rule. For participants in the high mismatch condition, 80% of the false problems violated the parity rule, and for participants in the low mismatch condition 20% of the false problems violated the rule. In this way, the base rate of success for using the parity-check strategy differed for the two groups of participants. This difference in base rate was reflected in the rejection response time advantage for parity mismatch problems, relative to parity match problems. The parity effect for participants in the high mismatch condition was almost twice that of participants in the low mismatch condition. This differential parity effect can be interpreted as participants being more likely to use the parity-check strategy when it had a higher base rate of success.

Lemaire and Reder's (1999) participants were not aware of either their use of the parity-check strategy or the percentage of parity mismatch problems. Only 5 of 32 participants reported using any strategy other than verification, and only 2 of those participants reported using the parity-check strategy. When participants were asked to guess what percentage of the problems were parity mismatch problems, their responses did not differ for the two conditions (high and low mismatch). Approximately half of the participants in each condition estimated that more than 50% of the incorrect problems violated parity, and the other half estimated that less than 50% violated parity. Thus, the participants did not seem to have explicit access to the relevant proportions that were influencing their strategy selection.

Lovett and Anderson (1996) found evidence that base rates of strategy success affect strategy selection regardless of whether participants' deliberately tried to use them. The task was Lovett's Building Sticks Task (BST) in which participants must build a stick of a target length by adding and subtracting sticks from an infinite resource of sticks of 3 different lengths. The lengths of the target and resource sticks vary from trial to trial. There are essentially two strategies that participants can use for solving the task. Neither of these strategies is inherently correlated with perceptual features (i.e., the lengths) of the various sticks, hence base rates of success with each strategy can be varied independently of whether or not perceptual problem features "suggest" one strategy over the other. Participants solved twenty blocks of BST problems. The base rate of success for each strategy was varied across groups of trials. For example, participants received 5 blocks of trials with one strategy working most of the time, then five blocks with that strategy working less frequently, and so on. Participants' strategy choices generally tracked the variable base rates of success of the alternative strategies.

As with the two previous studies, awareness was not critical for base rates to affect strategy selection. Sixteen participants claimed to have solved the task by looking at the lengths of the sticks, and twenty participants claimed to have solved the task by choosing the more successful strategy. However, when strategy selection was investigated as a function of which of these two strategies participants reported using, there was no difference between the two groups in the proportion of trials on which they selected the more successful strategy. Thus, participants' likelihood of selecting a successful strategy was not affected by whether or not they reported using base rates of success, even though the base rates influenced strategy selection.

4. STRATEGY ADAPTATION AND CONSCIOUS AWARENESS

We have established that people select among strategies, do so with variability, and select strategies adaptively in response to changes in intrinsic and extrinsic factors. The critical issue on which we now focus is whether or not participants adapt as the result of conscious monitoring. First consider whether adaptation to changing base rates of success is due to conscious monitoring. The evidence discussed earlier indicates that this is unlikely. In Reder (1987), Lemaire and Reder (1999), and Lovett and Anderson (1996) most participants were not aware of either base rates of success or changes in base rates of success. Many participants were not even aware of which strategies they had used to accomplish the task.

There are other tasks that show sensitivity to base rates without awareness. Participants in a study by Reder and Weber (1997) performed a spatial localization task. On each trial a target appeared in one of four spatial locations on the screen, and participants responded by pressing one of four buttons to indicate the location of the target. On most of the trials a distractor as well as a target was presented. The probability of a distractor was systematically varied across these four locations, such that 60% of the distractors occurred in one position, 30% in another, 10% in another, and no distractors appeared in the other position. Reder and Weber found that participants were slower to respond to the target as a function of distractor location. Participants were slowest when the distractor appeared in a location that rarely contained it, specifically the 30% or 10% position. However, when the distractor appeared in a frequent (i.e., 60%) location participants responded as fast as they did when no distractor was present. Although participants learned to ignore the distractor when it occurred in a frequent location, as indicated by their latency patterns, when participants were questioned at the end of the study about the distribution of distractors across the four locations, they were not aware of the different rates of presentation. Most participants reported that the distribution of distractors was even across the positions, the distribution was random, or they paid no attention to the distractors. In sum, participants showed sensitivity to the rates at which distractors were presented in various spatial locations without conscious awareness that there were different rates.

Chun and Jiang (1998) demonstrated that implicit learning of display contexts can help direct spatial attention to the location of a target. Their participants had to locate and identify a target in a display with several distractors. The context of each target was operationalized as the spatial layout of objects present on a trial. For half of the trials, the contextual configurations were yoked to target locations such that the target position could be predicted by the configuration of the other elements in the display. For the other trials, the configurations were unique and random. Participants detected targets in the repeated configurations more rapidly than targets presented in unique configurations. This benefit was maintained even when the actual distractor objects changed halfway through the experiment. Thus, participants' strategies for locating and responding to the target differed for the learned repeated configuration and the unique configurations. Importantly, this learning occurred without participants' conscious awareness of the repeated configurations. At the end of the experiment participants were given a recognition test of the configurations, and their hit rate and false alarm rate were the same (i.e., presented and not-presented configurations were indistinguishable). Only 3 of the 14 participants noticed that there were repeated configurations, and their hit rate was the same as their false alarm rate. Additionally, the latency benefit for repeated configurations did not differ for the aware and unaware participants. Clearly, conscious awareness of the repeated contexts, or even that there were repeated contexts, was not necessary for participants to learn about them or to have this knowledge affect performance (see Marescaux, Izaute, & Chambres, this volume, for a discussion of implicit learning and metaknowledge). This is yet another example of strategy selection occurring without conscious awareness of the factors that influence which strategy is used.

Sometimes the impression of conscious control of strategy choice may be illusory. In some tasks, participants may observe what they are doing and confabulate an explanation for their behavior that may or may not be accurate. Regardless of explanatory accuracy, when the rationale for strategy selection only comes after one's behavior it is not conscious deliberation that affects the regulating processes and strategies that are evoked.

Is performance better when participants are (or report being) aware of the factors that influence their strategy selections? The answer is "not necessarily." Recall Lovett and Anderson's (1996) study with the Building Sticks Task, in which participants' strategy selection was sensitive to base rates of success. Performance was no better for the participants who were aware of different base rates of success than for those who were not aware.

There are several factors that might affect whether or not a person is aware of the strategy that she has selected for performing a task. People are likely to be aware of the strategy selected when they are required to report it, as in the feeling of knowing and game show paradigm experiments previously discussed. When a particular strategy does not lead to a successful outcome, individuals may be inclined to assess the strategy that they are using in order to modify it or switch to a potentially more successful strategy. For example, a tennis player may wonder "Why are my tennis balls going into the net?" The timeframe of the unit task is also likely to influence people's awareness of their strategy selection. For processes that take on the order of a second to complete, there is no time for introspection to affect performance or strategy selection. There may be conscious strategy selection and metacognition for tasks that take several seconds (or longer) to complete, because there is more time to think about or reflect on what one is doing. The ease with which a task can be performed likely affects awareness, with easier tasks associated with less awareness. It may be that more difficult tasks elicit more awareness because they are more error prone or take longer to execute. Similarly, the degree of experience a person has with a particular type of task can influence his awareness of his strategy use. With practice the performance of some simple skills can become relatively automatic in nature (e.g., Shiffrin & Schneider, 1977), essentially eliminating the need for strategy selection.

5. CONCLUSION

This chapter has reviewed evidence that people's actions are influenced by features of the task, features of the environment, and their success at using specific strategies. Remarkably, this adaptation often occurs without any awareness of what procedures or strategies are being used, the base rates of types of stimuli in the environment, or the success of a given strategy. A number of these strategy choices, such as retrieving versus reasoning to the answer, have often been discussed as examples of behaviors that are subject to metacognitive control. Therefore, if people in the field maintain the position that metacognition requires conscious awareness, then we would argue that cognitive monitoring and strategy selection often occur without metacognitive intervention.

One way to respond to the evidence and claims that we present is to exclude a requirement of conscious awareness from the concept "metacognition" and distinguish between conscious and non-conscious (or explicit and implicit) metacognitive processes. Some researchers have recently adopted this view. For example, Koriat (2000) proposed that there are both information-based metacognitive judgments that involve conscious awareness and experience-based metacognitive judgments that are implicit in nature. Additionally, Sun (2000) argued that metacognition is not necessarily explicit nor implicit, but rather is a combination of both types of processes. For further discussion of the relationships between explicit knowledge, implicit learning and metacognition see Dienes and Perner (this volume).

Perhaps there are different degrees of metacognitive awareness that are involved in the control and monitoring of performance. For example, Carlson (1997) suggested that there are three levels of metacognitive awareness for problem solving (cf. Valot, this volume). He proposed that a problem solver's cognition can be reflective, deliberate, or routine. When problem solving is quite difficult, people may reflect on their problem-solving activities. The intermediate level refers to instances when a person considers alternative strategies and deliberately decides which strategy to try. In skilled problem solving, cognition is routine in that one does not deliberately select among strategies, but knows what to do and does it.

Whether or not one considers conscious awareness as a requirement of metacognition, it is important for people to know that cognitive monitoring and control often occur without awareness of either the factors that are influencing strategy selection or the strategies that are being used. Hence, when researchers find individual variability in strategy use they should not assume that strategy selection was the result of conscious deliberation.

Author note

Preparation of this chapter was partially supported by grants 2-R01-MH52808 and 5-T32-MH19983 from the National Institute of Mental Health.